SEcure Cloud computing for CRitical infrastructure IT

Contract No 312758

Deliverable: 3.1
Methodology for Risk Assessment and Management

AIT Austrian Institute of Technology • ETRA Investigación y Desarrollo •
Fraunhofer Institute for Experimental Software Engineering IESE •
Karlsruhe Institute of Technology • NEC Europe Ltd. • Lancaster University • Mirasys •
Hellenic Telecommunications Organization OTE • Ayuntamiento de Valencia • Amaris

Document Control Information	
Title	Deliverable 3.1 - Methodology for Risk Assessment and Management
Editor	Paul Smith (AIT)
Author(s)	Jerry Busby (ULANC), Lucie Langer (AIT), Marcus Schöller (NEC), Noor Shirazi (ULANC), and Paul Smith (AIT)
Classification	☐ Red – highly sensitive information, limited access for: ☐ Yellow – restricted limited access for: ☐ Green – restricted to consortium members ☒ White – Public
Internal reviewer(s)	Christian Jung (IESE), Reinhard Schwarz (IESE), Evangelos Sfakianakis (OTE)
Review Status	☐ Draft ☐ WP manager accepted ☒ Coordinator accepted
Requested deadline	2013/12/31

Versions			
Version	Date	Change	Comment/Editor
1	2013/12/05	Version for review	Jerry Busby (ULANC), Lucie Langer (AIT), Marcus Schöller (NEC), Noor Shirazi (ULANC), and Paul Smith (AIT)
2	2013/12/19	Incorporated comments from review	Paul Smith (AIT)
3	2013/12/28	Layout changes only	Roland Bless (KIT)

Abstract

Understanding the risks associated with using cloud computing is of paramount importance for operators of high-assurance ICT services, such as critical infrastructure providers. This is because failure of such services could result in disruptions to the core services that our society depends on, or the unauthorised disclosure of sensitive information. Furthermore, these organisations need to understand the risks, in order to properly assess the costs and benefits of cloud computing.

In this deliverable, we provide a number of items that can be used by organisations to understand these risks. A cloud-specific threat and vulnerability catalogue is presented, which can be used to support the implementation of a risk assessment. Furthermore, we show how this catalogue can be applied by organisations to understand the risks of cloud adoption, in the context of a novel extension to asset-driven risk assessment approaches. For a number of reasons, perceptions of risk are important to understand. We have conducted an analysis of individual and organisational perceptions of risk. The results of this analysis suggest that a major perceived threat to organisations are the intentions of the cloud provider, with respect to their data, and not necessarily threats emerging from cyber-attacks. We present initial findings on the features of a cloud infrastructure that can be measured to better understand risk factors, and observations regarding risk management for the cloud.

Contents

List of Abbreviations

Abbreviation	Expansion
AAA	Authentication, Authorisation and Accounting
ATCA	Advanced Telecommunications Computing Architecture
API	Application Programming Interface
CI	Critical Infrastructure
CIAMAU	Confidentiality, Integrity, Availability, Multi-party Trust, Auditability and Usability
CIP	Critical Infrastructure Protection
COTS	Commercial-off-the-shelf (system)
CPU	Central Processing Unit
CSA	Cloud Security Alliance
CSV	Comma Separated Values
(D)DoS	(Distributed) Denial of Service
DFD	Data Flow Diagram
DsAHP	Dempster-Shafer Analytical Hierarchy Process
E-DoS	Economic Denial of Service
ENISA	European Network and Information Security Agency
HW	Hardware
IaaS	Infrastructure as a Service
IT	Information Technology
ICT	Information and Communications Technology
IPSec	Internet Protocol Security
ISMS	Information Security Management System
ISO	International Organization for Standards
ISP	Internet Service Provider
ISRAM	Information Security Risk Analysis Method
KPI	Key Performance Indicator
MITM	Man-in-the-middle (attack)
LAN	Local Area Network
NERC	North American Electrical Reliability Corporation
NIS	Network and Information Security
NIST	National Institute of Standards and Technology
NSA	National Security Agency
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
OS	Operating System
PaaS	Platform as a Service
PDCA	Plan-Do-Check-Act
QoS	Quality of Service
QUIRC	Quantitative Impact and Risk Assessment Framework for Cloud Security
RAaaS	Risk Assessment as a Service
SaaS	Software as a Service
SLA	Service Level Agreement
SOA	Service Oriented Architecture
SQL	Structured Query Language
STRIDE	Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege
SCADA	Supervisory Control and Data Acquisition
URL	Uniform Resource Locator
VM	Virtual Machine
VMM	Virtual Machine Manager
VPN	Virtual Private Network

1 Introduction

The use of cloud computing has a number of potentially significant benefits for organisations that operate high-assurance ICT services. These organisations often provide the critical infrastructure services that our society depends on, such as gas and electricity networks, eGovernment services, and the ICT services that support the banking sector. The benefits of cloud computing include reduced capital expenditure on hardware equipment, the ability to more readily cope with fluctuating service demands, and increased fault-tolerance via the use of virtualisation, for example. However, the use of cloud computing introduces a number of new risks, which must be understood at the point when an organisation is considering adopting the cloud, and on a continuing basis. Understanding these risk is particularly important for organisations that operate high-assurance ICT services. In this deliverable, we support organisations that are considering or currently using cloud computing, in order to understand its risks.

There are a number of existing processes and guidelines that can be followed by an organisation to understand the cyber-security risks that could affect their ICT services. In many cases, an organisation in the high-assurance domain will be implementing one of these, in order to, for example, acquire ISO 27001 certification. Furthermore, with the proposed European Network and Information Security (NIS) directive [Eur13] potentially coming into force, the number of organisations carrying out such an information security risk assessment is likely to increase. We summarise a number of such risk assessment processes in Section 2, including those that are specifically related to industrial control systems.

Building on these relatively well-known risk assessment guidelines, we survey the related work that is focused on the issue of cyber-security risk assessment for the cloud. There has been a great deal of existing research in this area, including the identification and classification of the different threats and vulnerabilities that are associated with cloud computing. These are both technical and organisational in nature. Furthermore, a number of risk assessment processes have been proposed by the research community that target cloud computing – these often take the form of extensions or adjustments to existing approaches, in order to accommodate the specific challenges of risk assessment for cloud. These challenges include the lack of transparency about risk factors in multi-stakeholder deployment contexts and the dynamic nature of cloud infrastructures. This has lead the research community to propose the concept of Risk Assessment as a Service (RAaaS) – on-demand and online risk assessment, wherein properties of the cloud infrastructure are measured and modelled, in order to provide a continuous risk assessment. These issues and concepts are discussed further in Section 3.

Perceptions of risk in the context of cloud computing are important to understand, for a number of reasons. Such perceptions will inevitably influence decisions about cloud adoption, for example, and the nature of the security controls that will be applied. Furthermore, in many cases, risk assessments are augmented by, or primarily-based on, the views of experts, e.g., when using the Delphi method [LT75] for risk analysis. With this in mind, we have carried out a survey on risk perceptions for the cloud, on an individual and organisational basis. To garner individual perceptions of risk, we have implemented a survey questionnaire, which was disseminated to members of the SECCRIT user and advisory board, and a number of other individuals. Meanwhile, to understand the organisational perceptions of risks of cloud usage, we carried out an analysis of a university policy document, with respect to cloud computing. A key finding of this analysis is the primary risks the institution foresees (and introduces protection measures to address) are risks originating from the cloud provider via, what the institution perceives as, inappropriate use of their data. Further details of these analyses are discussed in Section 5.

Two important factors, amongst others, that must be considered when understanding cyber-security risks include: *(i)* the threats and their likelihood; and *(ii)* the vulnerabilities and an indication of their severity. A key challenge when understanding the risks associated with cloud computing is to determine those that are *specific* to the use of cloud. Without this understanding it is challenging to determine the additional risk (or otherwise) that emerges from the use of cloud. Therefore, to support an analyst to understand *cloud-specific* risks, we have developed a *cloud threat and vulnerability catalogue*, which is described in Section 6. As a basis for this catalogue, we have identified (based on previous research) a number of categories that enable us to focus directly on cloud-related issues – the core of these categories is based on the NIST essential cloud computing characteristics [MG11]. The catalogue was formed using two methods: a structured analysis of the existing literature, from an academic and industry perspective, and an analysis of the SECCRIT cloud architecture model, which is described in SECCRIT deliverable D5.1 [Rol13].

Central to carrying out a risk assessment is an understanding of what are an organisation's *security objectives*, which typically relate to its business processes and assets. This basis is required in order to understand how the various threats and vulnerabilities may affect an organisation. In many cases, these objectives are very much organisation-specific, in that they relate directly to their processes and assets. However, in Section 4, we discuss the potentially different forms of security *and resilience* properties that an organisation may be concerned with, and how they could be affected by the use of cloud computing. We discuss resilience objectives, in particular, because the organisations that operate high-assurance ICT services will very likely be concerned with these factors; this is in addition to those considered for classical security, such as confidentiality, integrity and availability.

Through engagement with high-assurance ICT service providers, we have learned that one of the largest difficulties they face, with respect to cloud computing, is determining the risk associated with adopting the cloud, versus continuing with their current deployment model. Or, stated more plainly, it is not clear to organisations whether the risks outweigh the potential benefits. To support organisations with this decision-making process, in Section 7, we define an extension to existing asset-driven risk assessment processes. In short, the process takes the results from an existing (non-cloud deployment) and augments it with the risks associated with cloud deployment by modelling a potential cloud offering, and augmenting the risk scenarios based on this model. The SECCRIT threat and vulnerability catalogue can be used to support this analysis. Furthermore, we show how the process can be applied through a video surveillance system scenario, using the Verinice Information Security Management System (ISMS) tool.

A significant problem with risk assessment, and with cloud-based risk assessment in particular, is determining meaningful values for the likelihood and severity of threats and vulnerabilities, respectively. As a starting point to address this issue, in Section 8 we outline a number of measurable properties of a cloud infrastructure that could be used to improve understanding about these factors. These properties form the basis of what could be measured as a RAaaS offering by a cloud provider. We will investigate the potential applicability of these measures in future work in the project.

Finally, in Section 9 we briefly introduce issues around risk management for the cloud. Further work in this area is required, for example, in providing specific guidance about the content and nature of Service Level Agreements (SLAs), incident-response management processes, and security controls that can be used to counteract the threats and vulnerabilities that are outlined in our catalogue.

2 Risk Assessment and Management Approaches

We present an overview of a number of existing standards and guidelines for critical infrastructure information security risk assessment and management. These standards and guidelines can form the basis of understanding the risks associated with cloud adoption, alongside the threat and vulnerability catalogue that is outlined in Section 6. More specifically, in Section 7, we augment an existing risk assessment process, which can be used to obtain ISO 27001 certification, to enable an organisation to understand the risks associated with cloud adoption.

The ISO 27000 series of standards relate to information security management, risk management and security controls. ISO 27001 "...provide a model for establishing, implementing, operating, monitoring, reviewing, maintaining, and improving an Information Security Management System."[1] The standards make use of the *Plan-Do-Check-Act* (PDCA) model to structure the improvement process. Building on this, ISO 27002 is a code of practice for information security, which outlines potential controls and control mechanisms that may be implemented according to the guidance provided within ISO 27001. ISO 27003 focuses on providing help and guidance for the implementation of an Information Security Management System (ISMS). The next in the series, ISO 27004, provides guidance on the development and use of measures and measurement for the assessment of the effectiveness of an implemented information security management system and controls, as specified in ISO 27001. ISO 27005 covers information security risk management. ISO 27006 offers guidelines for accreditation. In general, the ISO 27000 standards primarily focus on security policies and security management strategies. NIST SP800-53 [Joi13] is similar to the ISO 27000 standards, and lists and classifies security control requirements, from which each security process should extract a baseline. Furthermore, it encompasses the risk management process, specifying all activities for the selection of security controls to their application to an organisation's information systems. Meanwhile, ISO 31000 provides principles, a framework and a process for managing risk and aims at improving the identification of opportunities and threats, and effectively allocating and using resources for risk treatment[2]. ISO 31010 provides information on risk assessment concepts, processes and the selection of risk assessment techniques[3].

Behnia *et al.* [BRC12] point out the range of differences between security analysis methodologies, including OCTAVE [Sof13], ISRAM [KS05], CORAS [dBDG$^+$03] and several others. Furthermore, the European Network and Information Security Agency (ENISA) maintains a repository of risk assessment standards, methods and tools from a European perspective [ENI13]. The Operationally Critical Threat, Asset and Vulnerability Evaluation (OCTAVE) method [Sof13] can be used for identifying and managing information security risks. It contains methods, techniques and tools for an *asset-driven* evaluation approach, focusing on security practices and strategic issues, and self-direction. Similarly, Magerit is a risk analysis and management methodology that has been developed in Spain [CGC06]. In a similar manner to OCTAVE, it is driven by an analysis of the assets that are associated with an organisation.

The Magerit methodology has formed the basis for an information security management method that has been developed as part of the EU-funded PRECYSE project[4], which is focused on industrial control system security. Furthermore, specific to industrial control systems are the ISA99 [Int13], the NERC CIP [Nor13], and the NIST SP800-82 [SFS11] standards. ISA99 proposes an approach that starts from the confidentiality, integrity and availability objectives for information security, taking into account the specific priorities characterising automation

[1] http://www.iso.org/iso/home/standards/management-standards/iso27001.htm

[2] http://www.iso.org/iso/home/standards/iso31000.htm

[3] http://www.iso.org/iso/catalogue_detail?csnumber=51073

[4] The EU-funded PRECYSE project: https://www.precyse.eu

systems with respect to standard IT infrastructures. The NERC CIP standards are published by the North American Electric Reliability Corporation, and contain best practices for securing power systems. Finally, NIST SP800-82 provides best practices on architecture design and security controls for SCADA and industrial control systems. According to a recent ENISA study [LET+11], ISO 27002 for information security management is the most adopted standard, followed by NERC-CIP, NIST SP 800-82, and finally ISA99.

3 Related work

In this section, we introduce items of notable related work that address issues associated with conducting risk assessment for the cloud.

3.1 Cloud-oriented Risk Analysis

A number of cloud-specific risk assessment techniques and processes have been developed. For example, Saripalli and Walters propose the QUIRC framework – Quantitative Impact and Risk Assessment Framework for Cloud Security – for risk assessment in the cloud [SW10]. As part of the framework, a set of *security objectives* are defined, using the acronym CIAMAU: Confidentiality, Integrity, Availability, Multi-party Trust, Auditability and Usability. The latter three objectives (MAU) relate specifically to cloud issues. Multi-party trust is important in the cloud, as service deployment, e.g., in the context of a public or community provisioning model, requires trust between the various stakeholders, and can be compromised by cyber-attacks. In the QUIRC framework risks are determined by understanding the probability of a threat occurring – a catalogue is given of cloud-specific and web-based threats – and its impact on the aforementioned security objectives. This information is derived using Wideband Delphi analysis [LT75] – an expert-driven analysis approach. For our risk measurement catalogue described in Section 8, we similarly focus on security objectives and the risks to them.

Wang *et al.* propose to make use of *attack-defence* trees for carrying out a threat analysis in the context of clouds [WLK+12]. Attacks trees [Sch99] are a well-known deductive threat modelling approach that can be used to explore the potential ways an attacker can realise an attack goal. Defence trees can be used to augment an attack tree, in order to explore how protection measures hamper an attacker. Previous work has shown how attack-defence (or simply defence) trees can be used to support decision making when considering the most cost-effective protection measures to use [BDP06]. As we will discuss in Section 8, the major shortcoming of an approach such as this in the cloud relates to garnering sufficient knowledge about the various stakeholder's infrastructure, which may be commercially sensitive, in order to conduct the analysis.

The EU-funded OPTIMIS project[5], amongst other items, carried out research into risk assessment for the cloud, which is largely summarised in one of their deliverables [JKKD12]. In short, the project developed a risk assessment method that can be applied at different stages of the cloud service provisioning lifecycle – at deployment time and when the service is in operation. Risk can be assessed from the perspective of a *service provider* or an *infrastructure provider*. A number of functional components and a very high-level architecture have been identified to support risk assessment from the perspective of the two identified stakeholders. The outcomes from the risk assessment are used elsewhere in the broader OPTIMIS framework for the cloud, e.g., in the context of access control (allowing new services to be deployed).

[5] `http://www.optimis-project.eu/`

One aspect of the OPTIMIS risk assessment method addresses the issues associated with the uncertainty of assessing different infrastructure providers (at the point of potential deployment), caused by a lack of information provided by them. At the core of this is an approach called Dempster-Shafer Analytical Hierarchy Process (DsAHP) [HGX08] – a technique to support decision making using incomplete information about a number of criteria. The OPTIMIS project highlight the following criteria that can be used to evaluate cloud infrastructure providers:

- *Past Service Level Agreement (SLA) performance*
  This criteria relates to how well did the given cloud infrastructure provider meet its SLA targets – failures result in higher risk measures; conformance reduces the risk. It is not clear to which extent this, potentially commercially sensitive, information is published.
- *Geographical information*
  The location of an infrastructure provider's data centres and operations can affect risk in a number of ways, e.g., because of jurisdictional threats (see Section 3.3), political instability, or actual geographical phenomena, such as the propensity for earthquakes in a region or the availability of reliable power sources.
- *Certification and standards compliance*
  Compliance with standards requires the implementation of minimum best-practices within an organisation, e.g., associated with information security management processes that are outlined in the ISO 27000 series of standards. It is assumed that compliance leads to a reduced risk.
- *Business stability*
  The stability of a stakeholder organisation that form part of the cloud provisioning chain, from a business perspective, can reduce the likelihood of contractual threats emerging, for example, as discussed in Section 3.3.
- *General infrastructure, security, and privacy practice*
  These are the general practices that an organisation puts in place to manage infrastructure, security and privacy aspects. In many cases, these are arguably the items that will be challenging to assess as a third-party organisation in a multi-stakeholder cloud deployment setting.

The use of these criteria and the DsAHP approach allows a service provider to compare the risks associated with a number of cloud offerings, with incomplete information about each of the criteria. As mentioned earlier, perhaps the most challenging aspect is acquiring meaningful data about the internal practices of a third-party organisation.

Acknowledging the specific challenges of implementing risk assessment for the cloud, the concept of Risk Assessment as a Service (RAaaS), whereby a *continuous* and *on-demand* risk assessment can be performed, has been proposed by a number of researchers [TTG13, KP10]. Kaliski and Pauley suggest the following properties, which are aligned to NIST's essential cloud characteristics, of cloud computing that make it challenging to implement a risk assessment [KP10], and therefore motivate the need for RAaaS:

- *On-demand self-service*
  Automated on-demand service provisioning removes an important previously human operator implemented control point, e.g., related to verifying security and privacy settings. This leads to poor visibility with these aspects, thus hampering understanding of the risks.
- *Broad network access*
  Network access via the Internet, for example, significantly changes the attack surface

from being relatively static and understood, i.e., within an enterprise, to a dynamic set of endpoints with varying and potentially unknown security postures and capabilities.

- *Resource pooling*
  The impact on security of other tenants on a given set of services is not clear (see Section 3.3 for a discussion on threats related to resource pooling); also, a priori knowledge of resource availability is likely to be not forthcoming because of the varying demands of tenants.

- *Rapid elasticity*
  This property of cloud computing exacerbates the previous problems, and potentially extends the scope of a risk assessment from the (primary) cloud infrastructure provider and customer to other sub-providers, which are invoked when scaling-out.

- *Measured service*
  The systems that automatically collect metering data potentially represent a further vulnerability, thus introducing new risks that should be considered.

These issues are further compounded by the classic problem in information security risk assessment of obtaining reliable historical statistical data regarding threat likelihoods. Theoharidou *et al.* propose a number of requirements for RAaaS [TTG13]: The service should support the dynamic and continuous measurement of accurate, trusted and real-time data for a specific cloud deployment. These measurements should be made using both comprehensive *qualitative* and *quantitative* metrics, which are targeted to the cloud environment. Theoharidou *et al.* argue that RAaaS should be supported by a knowledge-base, e.g., in the form of an ontology, that collates the knowledge from publicly available resources, in order to address the lack of statistical data. Finally, modelling tools should be developed to support the analysis of threats in different attack and deployment scenarios. In Section 8, we define a number of metrics for risk assessment in the cloud, which represent initial work in the direction of realising such a risk assessment service.

3.2 Cloud Vulnerabilities

When performing a risk assessment it is necessary to understand the nature of the vulnerabilities associated with a target of evaluation, in this case, a cloud infrastructure. Grobauer *et al.* [GWS11] have explored cloud-specific vulnerabilities – an important aspect when considering the risks associated with moving services to the cloud, or remaining with a conventional deployment model. They approach this by initially defining what constitutes a *cloud-specific* vulnerability. In summary, it is suggested that a vulnerability is cloud-specific if it:

1. *“is intrinsic to or prevalent in a core cloud computing technology,*
2. *has its root cause in one of NIST’s essential cloud characteristics [MG11],*
3. *is caused when cloud innovations make tried-and-tested security controls difficult or impossible to implement, or*
4. *is prevalent in established state-of-the-art cloud offerings”*

Regarding the first item above, there are arguably three core cloud computing technologies: *(i)* web applications and services; *(ii)* virtualisation; and *(iii)* cryptography. Vulnerabilities associated with these could include hijacking or riding of web sessions, virtual machine escape, whereby a process acquires illegitimate access to resources outside the scope of its virtual machine, and obsolete or insecure cryptography, respectively. The vulnerabilities associated with NIST's cloud computing characteristics are similar to those identified by Kaliski and Pauley [KP10]. Furthermore, Grobauer *et al.* propose that challenges of implementing security controls due to the nature of cloud technologies, such as virtualisation, represent cloud-specific vulnerabilities. For example, in deployment models whereby tenants share an infrastructure, potential limitations on the ability of a security analyst to perform vulnerability scans of the infrastructure represent a vulnerability. Finally, vulnerabilities in state-of-the-art cloud service offerings could include so-called *injection vulnerabilities*, such as SQL injection attacks, and weak authentication – for example, contemporary cloud offerings typically make use of single-factor authentication, in order to access services. In our work, we build on this classification of vulnerabilities for the cloud, and further explore specific instances.

3.3 Cloud Threats

Molner and Schechter present a classification of cloud-related threats, not all of which relate to cyber-attacks, that are both *technical* and *non-technical* [MS10]. Of course technical threats, such as Denial of Service (DoS) attacks are of paramount importance. However, they suggest that non-technical threats are just as pernicious, and describe the following classes of threat:

1. *Contractual*
   These are threats that relate to contractual issues when using the cloud, such as bankruptcy of one of the stakeholders in the provisioning chain, potential switching costs between providers, and cost-overruns due to attacks that aim to maliciously consume resources that must be paid for. Whilst these threats do not always relate to cyber-attacks, they can nonetheless result in the reduced availability of a service.

2. *Jurisdictional*
   These threats relate to potential indirect legal coercion, e.g., cease and desist requests sent to the cloud infrastructure provider, and direct or indirect jurisdictional exposure from copyright holders and law enforcement agencies, such as exposure to "secret searches."

3. *Organisational*
   Threats of this nature can be challenging to assess, and relate to items such as human resources not being suitably screened and audited. An analyst can look to the application of organisational quality management standards, such as those related to the ISO 9000 standards[6], for determining whether these aspects are potentially being addressed.

Shared tenancy of a physical data centre infrastructure is one of the benefits of cloud computing, e.g., in public or community cloud offerings. However, shared tenancy introduces a number of potential new threats, which are explored by Molner and Schechter: Given potential restrictions on forensic capabilities, caused by unwilling tenants, a threat of *diminished audit, detection, and incident response capabilities* may occur. Another non-technical threat associated with shared tenancy relates to so-called *jurisdictional collateral damage* – this threat relates to situations wherein law enforcement agencies request the shutdown of a data centre because of miscreant behaviour of one of its tenants. Finally, specific technical threats from other tenants can occur, such as direct breaches, side-channel attacks, denial of resources (via miscreant API usage) or resource theft, for example. We return to this classification in Section 6, where we discuss cloud-specific threats in more detail.

[6] ISO 9000 – Quality management: http://www.iso.org/iso/iso_9000

In general, risks from the *insider threat* – e.g., from disgruntled (ex-)employees – are challenging to estimate, but are often cited as one of the most likely threats to manifest, and can have the highest impact. Increasingly, it is challenging for enterprises to determine the boundaries of their organisation, in order to examine this threat – this is especially the case for cloud deployment models that involve multiple stakeholders to realise a service. Abbadi *et al.* [ANM11] describe a *model* and *process* that can be used to identify the insider threats in the context of cloud computing. The process functions by identifying the components associated with a cloud service (a model of these is provided), including the potential actions that can be carried out on the components (at different cloud "layers") and the credentials required to realise these actions. Subsequently, the actors that have access to these credentials are identified, which leads to a systematic understanding of the insider threat. A shortcoming of this approach is that detailed information about the cloud infrastructure, the actors and their credentials are potentially required, which may not be available in all settings, e.g., when a public cloud offering is used. This is a problem with similar research that aims to develop models in a cloud context that describe infrastructure, stakeholders and their credentials, in order to explore security issues in relation to security objectives [BMP$^+$13].

In order to support the practical analysis of threats in the cloud, the Cloud Security Alliance (CSA) maintains a catalogue of top threats – currently, the catalogue describes nine threat types [CSA13]. The threats are intended to reflect expert opinion on the current threat landscape for cloud. The current list of top threats include: *(i)* data breaches; *(ii)* data loss; *(iii)* account hijacking; *(iv)* insecure APIs; *(v)* Denial of Service; *(vi)* malicious insiders; *(vii)* abuse of cloud services; *(viii)* insufficient due diligence; and *(ix)* shared technology issues. As a staring point, a risk analyst could explore the probability and potential impact of threats that manifest for each of these classes.

3.4 Summary

In summary, to the best of our knowledge, there have been two general forms of activity in the area of risk assessment for cloud computing: *(i)* developing understanding about the cloud-specific threats and vulnerabilities, including ways to organise these items into categories; and *(ii)* cloud-oriented risk assessment processes and methods. Regarding the development of understanding of cloud-specific vulnerabilities and threats, some useful work has been done, which is helpful for understanding risks in this context. However, threats and vulnerabilities are typically considered separately, i.e., there has not been one overarching way to consider them. In this deliverable, we bring these two complementary concerns together, and propose an overarching categorisation of threats and vulnerabilities, and systematically populate a catalogue.

A number of cloud-oriented risk assessment processes have been proposed. Many of these are based on existing techniques and processes, which are augmented for the cloud. For the most part, we would argue they are deficient in two major ways: *(i)* they do not provide specific guidance on the risk of adopting cloud versus remaining with an organisation's existing deployment model – a key concern for organisations that operate high assurance ICT services; and *(ii)* little attention is paid to some of the specific challenges of conducting a risk assessment for cloud-based services, including a lack of transparency between stakeholders and the dynamic nature of the environment. In this deliverable, we provide steps that extend an existing risk assessment, which can be used to determine the risks associated with using cloud computing. As mentioned earlier, one approach to addressing some of the challenges of risk assessment in the cloud is the notion of Risk Assessment as a Service (RAaaS). We suggest this concept has a great deal of potential, but is still very much in its infancy. In this deliverable, we present initial indicators that could be measured as part of a RAaaS offering, with respect to the threat and vulnerability catalogue that we have developed.

4 Security and Resilience Objectives

Before organisations can understand the risks from the cloud-specific threats and vulnerabilities outlined in Section 6, they must attempt to understand their *security and resilience objectives* – a target set of measurable security and resilience properties that could be affected by a threat or vulnerability. These objectives relate to the business processes and ICT assets, e.g., in terms of data and services, an organisation supports and implements. Sterbenz *et al.* have developed a useful arrangement of measurable security and resilience properties that can be considered [SHc$^+$10]. This arrangement is depicted in Figure 1, with the measurable resilience properties shown on the right-hand side of the figure.

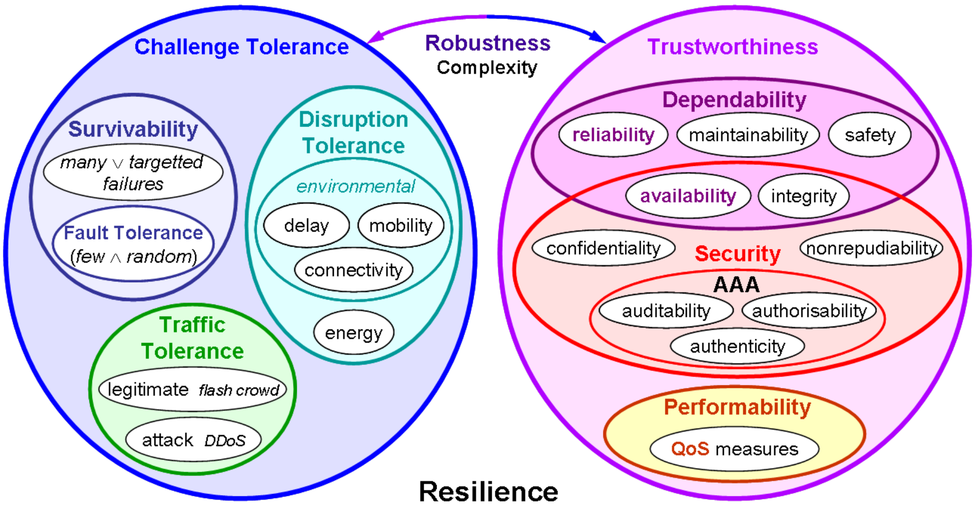

Figure 1: The disciplines that are related to network resilience, with the measurable objectives shown on the right

We briefly describe some of the pertinent security and resilience properties that are identified in Figure 1 (a more complete discussion on these properties can be found in the article by Sterbenz et al. [SHc$^+$10]), and relate them to deploying high assurance ICT services in the cloud:

- *Confidentiality* relates to the disclosure of information to unauthorised individuals or systems. As discussed in Section 6, there are a number of threats and vulnerabilities that are specific to the cloud, which could lead to confidentiality being compromised. Confidentiality is a key concern when considering eGovernment services, for example, whereby the private data of citizens is being handled. The SECCRIT project deliverable D2.2 on legal fundamentals [BP13] discusses issues associated with European data protection law, which (amongst other items) relates to issues of maintaining confidentiality. The deliverable outlines the various roles and obligations that exist with respect to European data protection law.
- *Integrity* relates to ensuring the accuracy and consistency of data over its lifetime. As discussed in Section 6, the integrity of data can be affected via the use of broad network access technologies or the failure of virtual service migration.

- *Availability* is the probability that a system will be operable when needed [SHc$^+$10], and for high assurance ICT services, especially for those that support critical infrastructures, is often the key objective. Cloud computing can be used to improve the availability of ICT services, via the use of virtualisation technology which can support tolerance to underlying hardware failures, for example. However, a number of cloud-specific threats are outlined in Section 6 that must be considered by a potential cloud user.
- *Auditability* relates to the aptitude of a system to be examined for correctness. In many cloud offerings that involve multiple stakeholders, e.g., public and community offerings, this objective can become compromised, e.g., because of a lack of transparency between organisations.
- *QoS measures* are performability properties, such as delay, bandwidth and jitter. Because of the rapid elasticity characteristic of cloud computing, wherein services can be migrated and scaled-out across an infrastructure, such quality of service measures could be compromised between components that constitute a service, leading to degraded performance.

Understanding the relative importance of these objectives with respect to an organisation's processes and ICT assets forms the basis of a risk assessment. We will show in Section 7.2 how an understanding of the core security objectives – confidentiality, integrity and availability – can be used to support an organisation as it tries to determine whether to adopt cloud computing for its ICT services.

5 Perceptions of Risk in Cloud Computing

Risk perception is an important branch of risk studies, for some obvious reasons. First, risk managers have to respond to what concerns individuals and organisations in society. A 'risk-based' view of what matters is often helpfully objective, but ultimately what matters to most governments and firms is what the public think. Second, expert risk assessments can often be wrong. Sometimes this is through the limitations of scientific knowledge, and sometimes because experts do not fully understand the conditions in which risks arise. Non-expert individuals, who may know a lot more about the particular context in which they operate, could have a better idea of the specific risk they face than is conveyed by a generic expert assessment. Third, the risk is often influenced by the behaviour of the individual risk bearers, i.e., those responsible for managing the risk, and their risk perceptions often matter in shaping this behaviour.

A similar study to ours on risk perception was conducted by the EU-funded T-Clouds project [BHC13]. Given the different approaches to conducting the study, it is not readily possible to draw direct comparisons between our results. The survey conducted by the T-Clouds project suggests that their stakeholder's highest risk concerns include "cloud specific attacks by externals, accidental leakage of data and credentials, insider attacks (e.g. by cloud administrators), and insufficient protection against more general IT security risks and attacks." These items are captured in our threat and vulnerability catalogue, which is presented in Section 6.

5.1 Cloud Computing Risk Questionnaire

To get a better understanding of the perceptions of risk for cloud computing, we created a questionnaire. The aim of this questionnaire was to determine the risk perceptions of members of organisations that were either users or providers of cloud computing. The design of our survey was based on:

- A distinction drawn between real-time risks and organisation-evolution risks. We were interested in direct, short-term threats such as loss of confidentiality, integrity and availability. But we were also interested in the longer-term risks such as losing expertise that might be needed in the future, and losing ultimate control of how services develop.
- A concern both with the degree of risk and the way in which people are held responsible for it. Risks are conventionally defined in terms of probability and impact, so we asked about these for various categories of risk. But creating a risk also creates responsibility for controlling or managing it.

There were three parts to the questionnaire:

1. The first asked for information about the respondent and the organisation they belonged to, including their views of what the main risks are and what kind of formal assessment such risks get in the organisation.
2. The second asked for assessments for stated categories of short-term risks.
3. The third asked for assessments for stated categories of organisation-evolution risks.

The analysis of the questionnaire was to be largely descriptive, intending to find out what mattered to organisations using and providing cloud services, and to find out whether differences between organisations could be explained in a systematic way. One problem with managing risk generally is that one party's view of a risk, and its decisions about how to deal with it, can increase the risk to another party. It therefore becomes important to understand how different parties see risks and the way they are managed.

5.1.1 Questionnaire Method

The questionnaire was administered online. All members of the advisory board were asked to complete it, and the researcher asked a range of personal contacts also to complete it. There were 17 responses of some kind, and different items received different numbers of responses, as indicated in the results that follow. The analysis was purely descriptive. An outline of the questionnaire is attached as an appendix.

5.1.2 Questionnaire Results

In terms of the demographics, Figure 2 shows the relative numbers of users and providers of cloud services, and the numbers in different industries. Figure 3 shows the length of (3(a)) personal then (3(b)) organisational experience in the main categories of cloud service, namely IaaS, PaaS and SaaS.

Respondents were asked, as users, what applications they used as cloud services. There were eleven responses, and applications included (in the respondents' descriptions):

- Email, calendars, contact databases
- Data storage

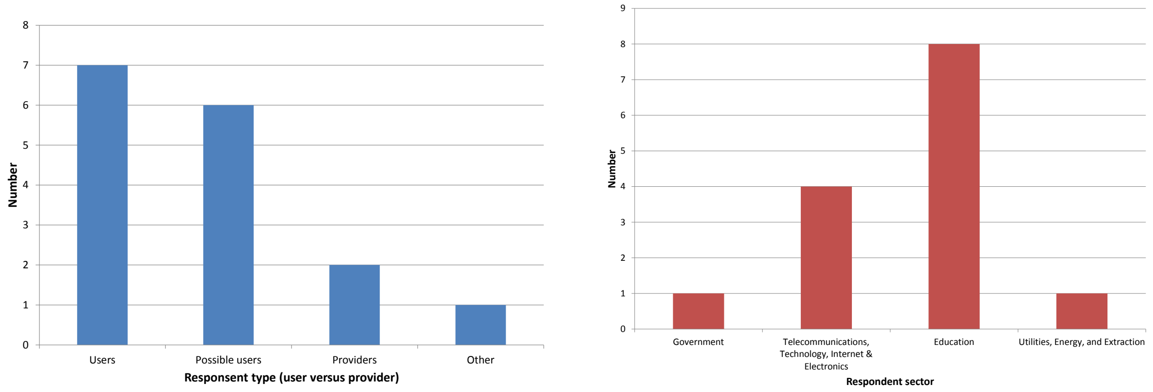

(a) Breakdown of users versus providers of cloud services

(b) Respondent sectors

Figure 2: Demographics of the respondents of the SECCRIT risk perceptions questionnaire

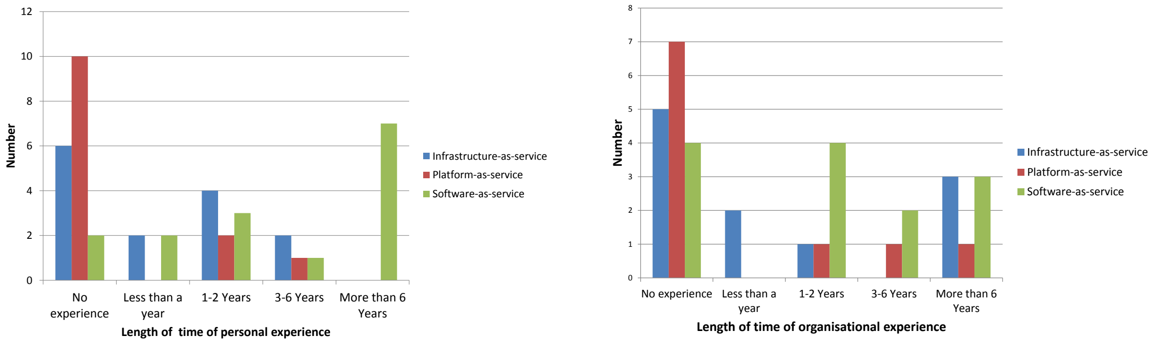

(a) Personal experience

(b) Organisational experience

Figure 3: The amount of experience the questionnaire respondents have with cloud computing

- Security services
- Social media
- Online payments
- Network planning tools
- Research applications

Respondents were asked what they thought were the main vulnerabilities of cloud services, in an open question, before being asked to rate risks in given categories. There were twelve responses (one of which was ‘no idea’), which are tabulated verbatim in Table 1. The first and tenth entries are from service providers, the others are from users or potential users. The responses point to the wide range of concerns that are associated with cloud computing, and considerable differences between the concerns of different individuals in different organisations.

Respondents were then asked, again in an open question, what risk assessment procedures their organisations applied to cloud computing. Again, it is worth quoting the non-null responses verbatim. The first and seventh entries, shown in Table 2, are from service providers, the remainder from users. Null responses presumably indicated that respondents did not know what risk assessments were being used, although they might also have indicated that there were no cloud-specific procedures. Again there is considerable variety in the responses.

Table 1: Perceptions of vulnerabilities in cloud services

No.	Vulnerability
1	Denial of Service / Outage; Data Confidentiality; Various application exploits due to limited security patching / endpoint protection; Unauthorised access (including failures in privilege escalation or internal attacks)
2	Delays due to Internet connectivity bandwidth variations.
3	Virtual machine escape; Session riding and hijacking; Insecure or obsolete cryptography.
4	Common known protocols and systems; Physical intrusion outside the own scope
5	Vulnerabilities exists in all technologies which are used in cloud computing such as virtualisation, Web 2.0 and SOA
6	Data Storage infrastructure; Network servers; storage servers; storage locations
7	Jurisdiction; Ownership of records and data; Availability of metadata; Audit capabilities; Rigidity of SLAs; Security; Preservation (as opposed to storage), retention and disposition
8	Data protection; User authentication
9	Confidentiality issues
10	Forms of denial of service that impact KPIs such as response latency
11	Failure; Denial of service; Privacy violations
12	The lack of a reporting structure for Cloud Services to tell users that their data may be at risk. The more popular a service is the more likely it will be targeted.

Table 2: Risk assessment procedures used by respondents organisations that are applied to cloud computing

No.	Procedures
1	Corporate Cloud Security Standard that is compliant with a mix of generic ISO and other vertical / mission-specific security or information assurance standards
2	In addition to the ISO standards we assigned an agency with intrusion detection
3	We do risk assessment of our devices in organisation
4	I am not aware of anything systematic
5	Calpana Crisam (http://www.calpana.com/)
6	Our concerns are currently blocked by legislative restrictions on where we can store our confidential data (neatly all our data potentially contains content covered by privacy laws)
7	We ask providers to abide to our internal best practice requirements
8	SLA
9	A qualitative risk assessment is used to identify if user data or authentication methods are exposed to attack. A statement of the security measures available at the Cloud Service provider is requested

In terms of the respondents' risk judgments, they were asked to rate both probability and impact of several categories of direct risk, using a number of options: unsure, very low (1), low (2), moderate (3), high (4), and very high (5). Probability and impact were multiplied to get a risk rating, and Figure 4 shows the mean rating, plus or minus one standard deviation for each category. The dominant risks are those of denial of service and threats of 'social engineering', and the risk of least concern was risk arising from having one user's data being co-located with another, highly targeted user.

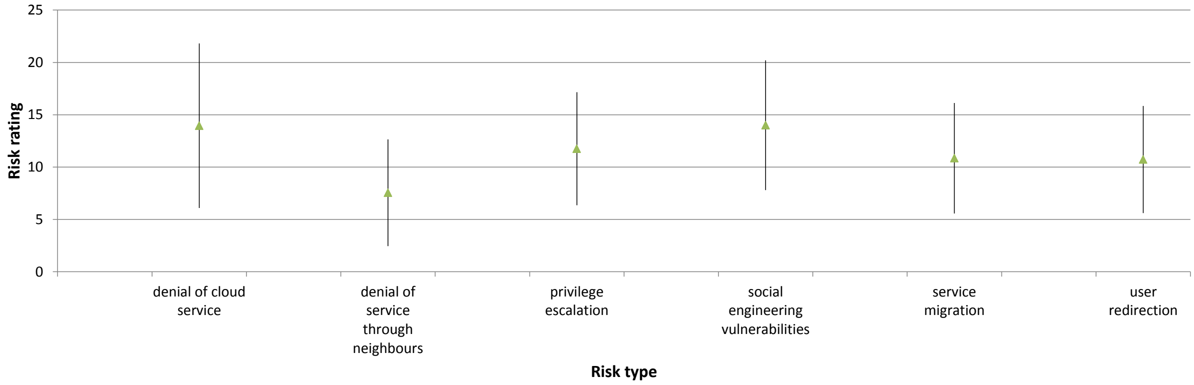

Figure 4: Judgment of direct risks by respondents, which are associated with cloud computing

There were only two respondents who claimed to belong to a service provider, so it is not possible to test whether users' and providers' judgments differed significantly. Looking at the response of the two providers, moreover, shows that they made very different judgments. Figure 5 shows their probability and consequence ratings for the different risk categories.

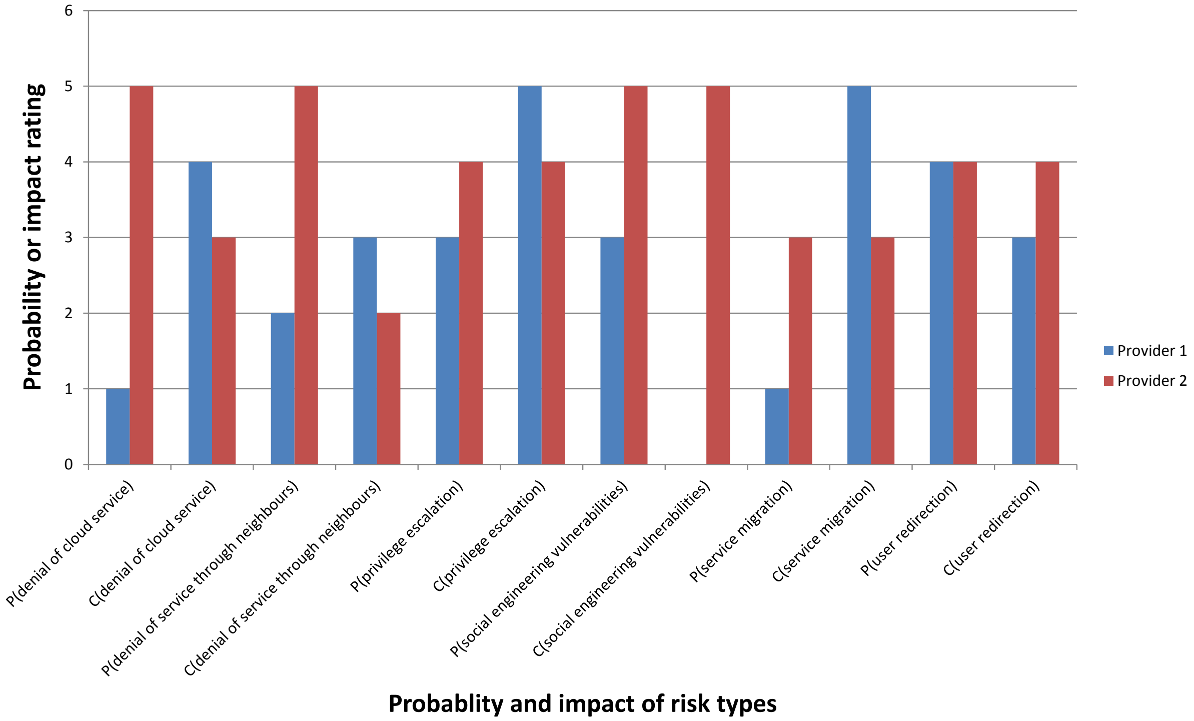

Figure 5: Cloud service providers risk probabilities and consequences for cloud threats

All respondents were asked to say whether each of several actors should take responsibility for these risks. Assigning a 1 to an actor for responsibility, and 0 for no responsibility, Figure 7 shows the total mass of responsibility attributed to each actor for each risk. Service providers carry the most attributed responsibility: the role of users and their system managers depends very much on the risk. Overall responsibilities are shown in Figure 6:

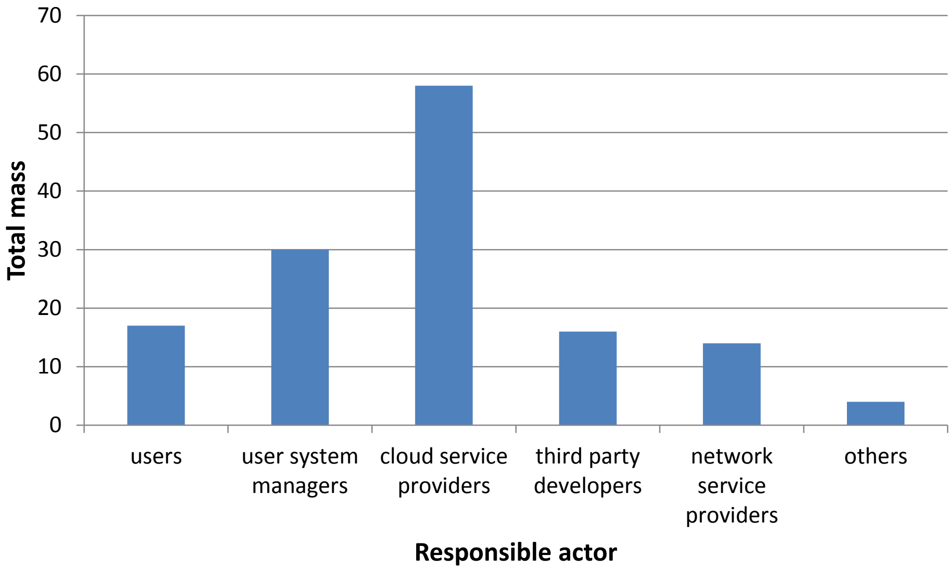

Figure 6: Overall actor responsibilities for risks

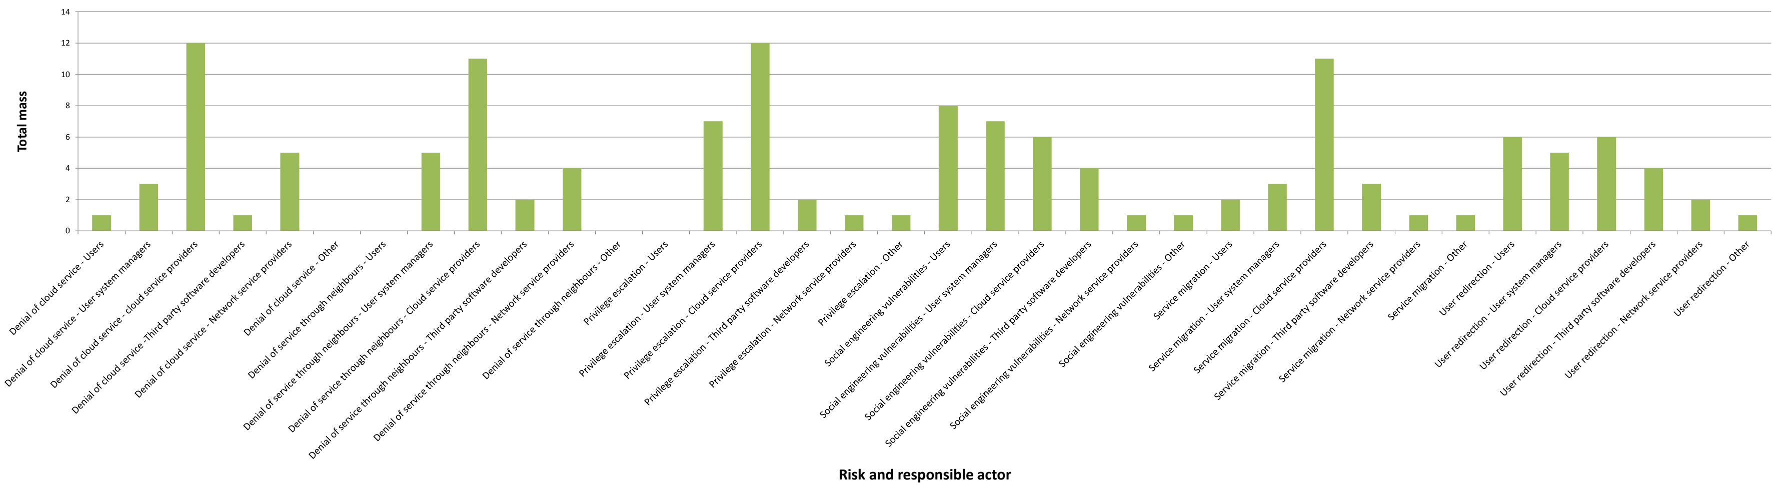

Figure 7: Respondents opinions on the actors that are responsible for specific risks

Respondents were also asked to rate risks categorised according to the element of the system at risk. These were rated on dimensions of vulnerability and impact. Figure 8 shows the product of the two, as mean plus or minus one standard deviation. The riskiest elements of the system in respondents' aggregate judgments were application software and the cloud service provider organisation, and network services the least risky, but there is not a great deal of variation.

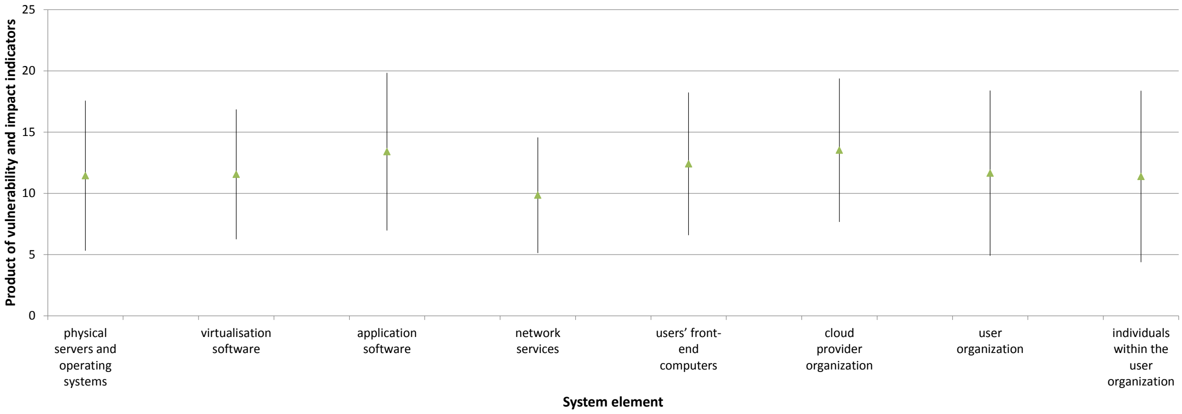

Figure 8: Measure of risk – the product of vulnerability and impact indicators – for different system elements

Longer-term risks, like direct risks, were rated on two dimensions – probability and impact. Again, we can use their product as an index of risk. Figure 9 shows risk ratings, in terms of mean plus or minus one standard deviation. The long-term loss of in-house IT capability was the risk of least concern, whereas the difficulty of tracing security failures was of the highest.

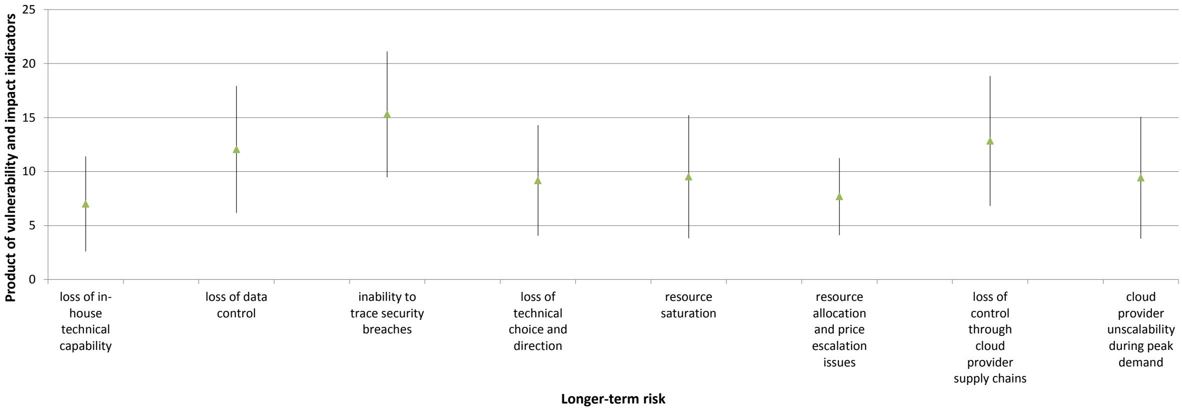

Figure 9: Longer-term risks – the product of vulnerability and impact indicators – emerging from cloud computing use

Three respondents also provided responses under the 'other' category. One in particular identified other long-term risks as:

- "Loss/fragmentation of IT security operations including assurance, compliance and security patch management (Likelihood: 4, Impact: 4)"
- "Loss of control over IT governance (Likelihood: 2, impact 4)"
- "Cloud provider lock-in (Likelihood: 4, Impact, 4)"

Another stated “Data confidentiality, privacy, issue of Trust, access, regulations, compliance.” And the third cited “Foreign legislation requiring access to our systems,” and referred specifically to the National Security Agency (NSA). One respondent also pointed out that the answers to some of our questions depended on who was providing the cloud service, particularly whether they were large-scale providers or small ones. Averages could not be taken as being representative of either.

5.2 Analysis of Organisational Policies as Expressions of Organisation Risk Perception

The need to work out what risks matter to people, how they respond to them, whether they agree with one another and so on, makes it important to study risk perception. But, as Tierney (1999) [Tie99] points out, we have to focus on organisations rather than individuals in order to understand risk and how it is responded to. Therefore, accompanying the survey on individual responses on risk associated with cloud computing, it is necessary to look at organisational responses. Perhaps the best way of doing this is to analyse the policies or other formal frameworks that organisations construct. These are, to varying degrees perhaps, deliberately constructed, collective responses that have enough credibility within an organisation to be published in some form. They should reflect the kind of risks that have organisational significance, and then describe the kinds of things that an organisation thinks it can do to deal with the risks.

5.2.1 Organisational Policy Analysis Method

Within the context of the SECCRIT project, we have only been able to demonstrate the kind of analysis that this would involve. Our analysis is based on the published cloud computing services audit of one particular organisation, a university. The content of the audit is likely to be very different from those in a commercial or governmental user of cloud services, or from those in a cloud service provider.

The method is essentially that of qualitative, grounded analysis. A text is read, categories are produced in the course of reading that – for the reader – capture what is essential to the meaning of the text, and the analysis then proceeds towards more abstract themes to generate some kind of theory. There is a process of ‘constant comparison’ through which parts of the text, and the categories used to organise them, are continually revisited in the process of finding theory. The aim is usually to avoid preconceptions as far as possible. In this case, there is a rather obvious notion that policies will *1)* say things that help us understand what risks have been noticed and become significant to the organisation; and *2)* say things that express what can and should be done about such risks.

5.2.2 Analysis Results

The immediate product of the analysis is shown in Table 3, containing a series of fragments from the document. Each is categorised either as identifying a risk, or identifying a risk control. Types of risk and types of controls are then distinguished. Some general points can then be inferred from this:

1. Risks are mostly related to concerns with the provider’s current and future intentions.

2. Only one risk concerns security of the service, with respect to only one specific service, and the text provides two URLs but does not expand on this risk – saying merely that it is 'worthy of note'.

3. The only formal control device is categorisation of data, with rule saying whether each category can be stored or transmitted via a cloud provider.

4. Informal controls consist of advice on redundant storage, encryption action, and ad hoc contracts.

5. Specific cloud services are itemised, indicating that the organisations expects to assess risk separately for different services, but the assessments are mostly similar, and there are statements of general concern such as "... no cloud-based service has been identified which is appropriate for Personal or Restricted data..."

5.2.3 Organisational Analysis Conclusions

The results broadly speak for themselves. This organisation is more concerned with risk originating in the service provider than in third parties attacking the service provider's facilities. And it mainly deals with this risk by setting simple organisational rules specifying what kind of information cloud services can be applied to. Repeating this simple analysis across multiple organisations would have several benefits:

1. It would help the organisations to learn from their peers (for example, other users) what kinds of risk others have identified, and what controls they have identified.

2. It would help us look for discrepancies between service providers and service users that seem likely either to constrain levels of trust below what is reasonable, or indicate levels of trust that are above what is reasonable.

3. It would help us look for inconsistencies in risk identification or risk control that might be problematic. It might point to unnecessary controls, to contradictory controls or to mutually supporting redundant controls (where one organisation's risk controls – for example, restrictions on its legal liability – create risks for the other organisation which has to implement controls in response).

Table 3: Extracts from the policy document used as part of the organisational risk perception analysis

Section	Subsection	Fragment	Main category	Sub category	Sub category
Introduction		"Anecdotally, the use of cloud based services is widespread at the institution"			
		"people are concerned and eager to have direction on the safe use of such services"			
		"the findings of an initial audit (in July 2012) of a few of the cloud based services known to be in use"			
iCloud	Issues	"Terms and Conditions, or nature of service can change without notice"	Risk	Risk to service terms	
	Issues	"No guarantee that data stays within the EU and US"	Risk	Risk to data location	
	Issues	"The right to change data in transmission is reserved"	Risk	Risk to data stability	
	Usage	"Personal Data: Not permitted ; Restricted: Not recommended ; Confidential: Permitted ; Ordinary: Permitted"	Control		Permissive rule
	Usage	"it is not always clear when the service is being used. For example, opening an attachment to a piece of mail within 'Pages' will implicitly move it to iCloud if Pages synchronisation is activated"	Risk	Risk to usage	
Gmail	Issues	"...When you upload or otherwise submit content to our Services, you give Google (and those we work with) a worldwide license to use, host, store, reproduce, modify, create derivative works (such as those resulting from translations, adaptations or other changes we make so that your content works better with our Services), communicate, publish, publicly perform, publicly display and distribute such content..."	Risk	Risk to data privacy	
	Issues	"Terms and Conditions, or nature of service can change without notice"	Risk	Risk to service terms	
	Issues	"No reference to where data is stored, hence no guarantee it remains in EU or under safe harbour"	Risk	Risk to data location	
	Usage	"Personal Data: Not permitted ; Restricted: Not recommended ; Confidential: Permitted ; Ordinary: Permitted"	Control		Permissive rule
	Usage	"Although personal data must not be transmitted via email within the institution, automatic forwarding to Google mail would exacerbate any break of that policy"	Risk	Risk to usage transparency	

Section	Subsection	Fragment	Main category	Sub category	Sub category
Evernote	Issues	“Terms and Conditions, or nature of service can change without notice”	Risk	Risk to service terms	
	Issues	“No guarantee that data stays within the EU and US”	Risk	Risk to data location	
	Issues	“A superseding agreement can be negotiated”	Risk	Risk to service terms	
	Usage	“Personal Data: Not permitted ; Restricted: Not recommended ; Confidential: Permitted; Ordinary: Permitted”	Control		Permissive rule
	Usage	“It is easier to identify when Evernote is being used than it is with iCloud, but caution is still recommended to ensure that personal data is not inadvertently moved into an Evernote account”	Risk	Risk to usage transparency	
	Usage	“If the institution wanted to pursue the use of a particular cloud-based service for use with Personal and Restricted data along with Confidential and Ordinary, it may be possible to draft an agreement with Evernote for a superseding agreement which ensures that data stays in the EU”	Control		Ad-hoc contract
Dropbox	Issues	“. . . ‘If you are using the Services on behalf of an organisation, you are agreeing to these Terms for that organisation and promising that you have the authority to bind that organisation to these terms’. . . ”	Risk	Risk to personal liability	
	Issues	“. . . ‘We may also remove any content from our Services at our discretion’. . . ”	Risk	Risk to service terms	
	Issues	“Terms and Conditions, or nature of service can change without notice.”	Risk	Risk to service terms	
	Issues	“No guarantee that data stays within the EU and US, though does claim to adhere to the US Safe Harbor laws.”	Risk	Risk to data location	
	Usage	“Personal Data: Not permitted* ; Restricted: Not recommended ; Confidential: Permitted ; Ordinary: Permitted”	Control		Permissive rule
	Usage	“it is easier than with some other cloud based services to apply further levels of security to files and still make them accessible via Dropbox. Hence, if there were unavoidable circumstances in which personal data needed to be stored on Dropbox, that, though not recommended, may be permitted as long as the file containing the data is separately encrypted”	Control		Encryption action
	Usage	“There have been security issues with Dropbox, which are worthy of note [cites URLs]”	Risk	Risk to data security	

Section	Subsection	Fragment	Main category	Sub category	Sub category
Advice		“Do not use cloud-based services to hold personal data unless it is independently encrypted, using a recommended encryption mechanism”	Control		Encryption action
		“Do not assume that your right to the Intellectual Property of documents held in the cloud is unchanged by storing the information there – many claims are made that you forgo certain rights through your use of certain cloud services”	Risk	Risk to intellectual property ownership	
		“As with advice on the use of mobile devices, do not rely on cloud-based services to provide more than a secondary copy of data; ensure that primary copies are held on University systems that offer suitable backup provision”	Control		Redundant storage
Conclusion		“whether that usage represents risk to information security is dependent on the category of information being shared”			
		“may be difficult to ensure that audit results remain up-to-date”	Meta-risk		
		“Dropbox’s claim that someone in an organisation accepts liability for the use of the service on behalf of the organisation; this requirement challenges the University to have an explicit policy about the use of their service”			
		“To date, no cloud-based service has been identified which is appropriate for Personal or Restricted data (based on their public terms and conditions)”		Control	Permissive rule
		“especially in the cases where it is not obvious when a cloud based service is being used (e.g. iCloud and Google mail with automatic forwarding)”	Risk		Risk to usage transparency

6 SECCRIT Cloud-related Vulnerabilities and Threats

In this section, we describe a catalogue of vulnerabilities and threats that apply when deploying high-assurance ICT services in the cloud. These can be used as a basis for carrying out a risk assessment. To compile this catalogue, two core activities were carried out: *(i)* a structured analysis of related material on cloud computing threats and vulnerabilities; and *(ii)* an architectural analysis, based on the SECCRIT architectural model. Figure 10 shows the management view of this architecture, which illustrates the interfaces used for deployment and management. It is necessary to point out that this view shows logical components and the interfaces between them only. In a commercial deployment, central components like the *Tenant Infrastructure Management System* are expected to be highly available by using suitable redundancy mechanisms. Thereby, single points of failure are avoided and the risk of them failing is deemed negligible. An in-depth discussion of the SECCRIT architecture can be found in Deliverable D5.1 [Rol13].

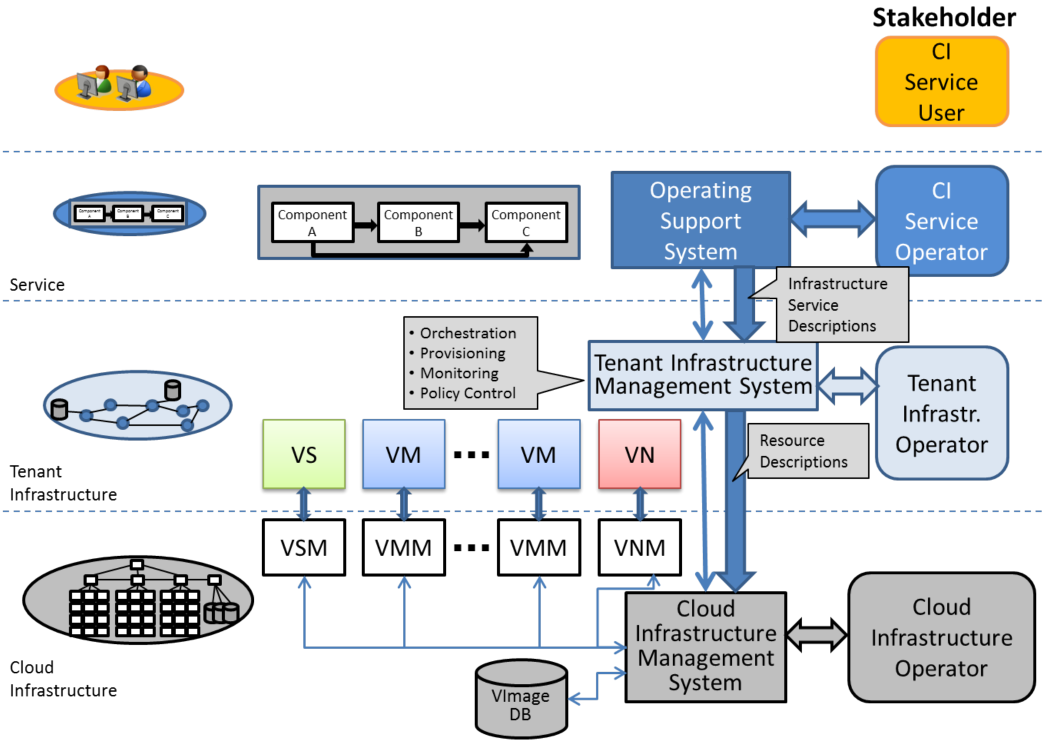

Figure 10: The SECCRIT cloud management architecture for critical infrastructure services

We have analysed this architecture for threats and vulnerabilities that do not exist in deployments using a *box-model*, i.e., when a service is run in a purpose-built environment, and not in the cloud. Furthermore, we analysed the differences in challenges to the operation of critical infrastructure services running in cloud environments. For the latter, we first had to compare the various options of how services can be run in cloud environments and compared them with the box-model. The box-model mode of operation is shown in Figure 11(a). Again, the figure shows a logical view; the service usually consists of a set of components cooperatively providing the specified service – these components are represented by the green box label “CI Service.” The simplest approach to virtualising such a service is to take the existing software, install it into a virtual machine (VM) image and execute it on top of a virtual machine manager (VMM). The VMM provides an abstraction from the underlying hardware, which we just call “compute” – the compute hardware can be any computing platform: Intel’s x86 platform and ARM platforms are dominant in data centre environments. Figure 11(b) shows this simple approach to virtualisation. From the critical infrastructure service point of view, the VM, VMM, and compute node constitute the platform it is running on.

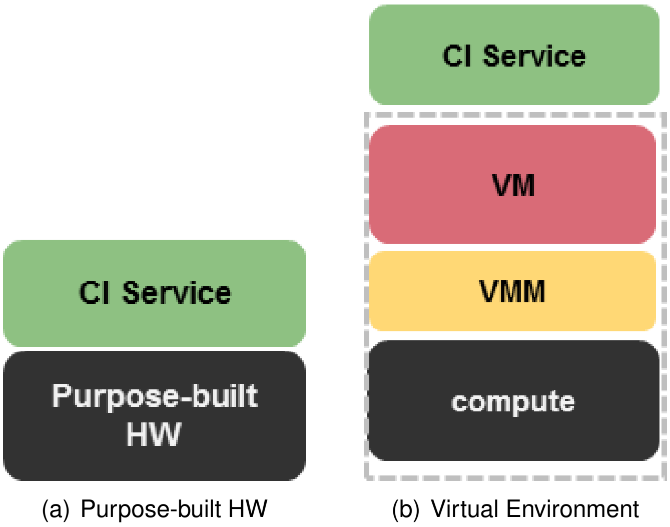

(a) Purpose-built HW (b) Virtual Environment

Figure 11: A critical infrastructure service running on purpose-built hardware versus within a virtualised environment

6.1 The SECCRIT Vulnerability and Threat Catalogue

As shown in Figure 12, we have organised the catalogue in a similar manner to Groubauer *et al.* [GWS11], by arranging them primarily into the five NIST essential cloud characteristics [MG11]: *(i)* on-demand self-service; *(ii)* broad network access; *(iii)* resource pooling; *(iv)* rapid elasticity; and *(v)* measured service [MG11]. Furthermore, we include categories that relate specifically to *virtualisation* as a key enabling technology, important *organisational issues*, the underlying *physical cloud infrastructure*, security and resilience *control implementation* challenges, and issues associated with contemporary *cloud offerings*. For each of the threats and vulnerabilities, we highlight the primary security and dependability objectives they affect – confidentiality (C), integrity (I) and availability (A) – they affect and, when possible, point to related material that can be used to obtain further details.

6.1.1 On-demand Self-service (ODSS)

The ability to automatically provision resources without the need for human-operator intervention – so-called on-demand self-service – is one of the fundamental characteristics of cloud computing [MG11]. Typically, this provisioning is undertaken via a management interface by the cloud user. Whilst this introduces a greater degree of flexibility and convenience, a number of security vulnerabilities and threats consequently emerge. For instance, Kaliski and Pauley [KP10] highlight this form of automation removes an important, previously human-operator-implemented, control point, e.g., related to verifying security and privacy settings, thus leading to poor visibility regarding these aspects (ODSS-1).

Furthermore, the use of a remote management interface to implement on-demand self-service introduces a number of vulnerabilities and threats that would otherwise not exist. Grobauer *et*

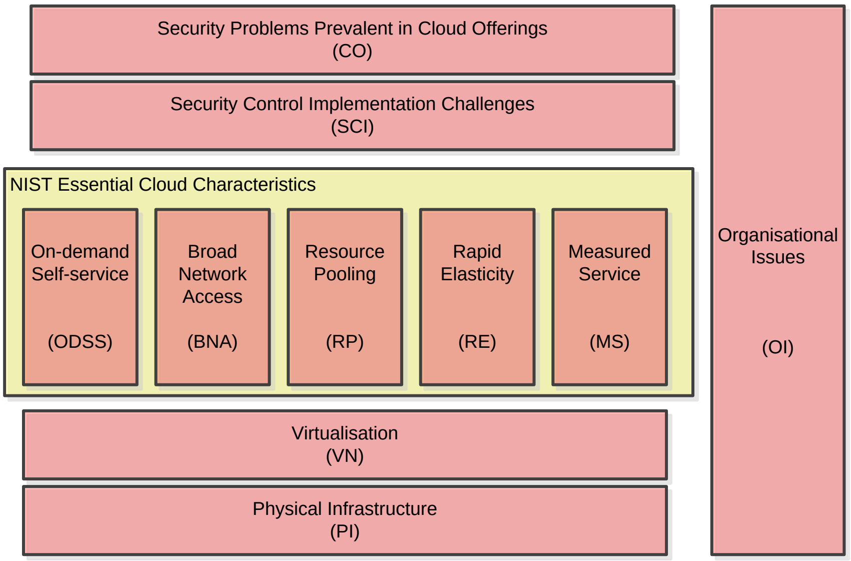
Security Problems Prevalent in Cloud Offerings
(CO)
Security Control Implementation Challenges
(SCI)
NIST Essential Cloud Characteristics
On-demand Self-service (ODSS)
Broad Network Access (BNA)
Resource Pooling (RP)
Rapid Elasticity (RE)
Measured Service (MS)
Virtualisation
(VN)
Physical Infrastructure
(PI)
Organisational Issues
(OI)

Figure 12: The SECCRIT threat and vulnerability catalogue categories

Table 4: On-demand Self-service-related vulnerabilities and threats

On-demand Self-service				
ID	Description	Type	Obj	Refs
ODSS-1	Loss of human-operated control point to verify security and privacy settings	Vu	C-I-A	[KP10]
ODSS-2	Poor authentication, e.g., single-factor authentication, on the cloud management interface	Vu	C-I-A	[GWS11]
ODSS-3	Denial of Service (DoS) attack against the the cloud management interface	Th	A	[MR04, CXZB11]
ODSS-4	Failure of the cloud management interface	Th	A	
ODSS-5	Man-in-the-middle attack on the management interface	Th	I	[CCR09]
ODSS-6	Uncontrolled request for resources	Th	I	

al. [GWS11], for example, highlight that in many cloud offerings poor, i.e., single-factor, authentication techniques are used to access such interfaces (ODSS-2). Clearly, the failure of the management interface for various reasons, including misconfigurations or software faults, represent a threat if the interface is required and proves to be unavailable (ODSS-3). This problem could be particularly pernicious if resources are needed to be provisioned via the interface in response to an unusual increase in service demand. Similarly, the management interface can be susceptible to a number of threats by malicious actors, such as (Distributed) Denial of Service (DoS) attacks (ODSS-4) [MR04, CXZB11], which can result in a lack of availability of the interface, and man-in-the-middle (MITM) attacks (ODSS-5) [CCR09], which can result in a cloud service provider's user account being hijacked, for example. Furthermore, full automation of the service deployment and life-cycle management system poses the risk of high costs if the interface to the virtual resource management malfunctions due to a software bug and requests resources in an uncontrolled manner. Specific care should be taken if the management interface is accessed via public networks, as we will discuss next.

6.1.2 Broad Network Access (BNA)

The characteristic “broad network access” relates to the capability that cloud services can be accessed via standard network mechanisms, such as the Internet protocol suite, that promote the use of a range of end-user device types [MG11]. From a deployment perspective, this can involve the use of the public Internet, or some other form of wide-area network. This is a departure from many non-cloud-based deployments, whereby services are typically accessed via protected local-area networks, and in some cases – particularly in the critical infrastructure sector – using proprietary or specialised (application-level) protocols.

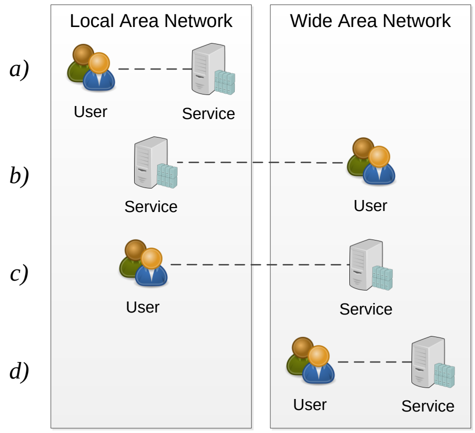

Figure 13: Potential network arrangements when deploying remote services

The use of wide-area networks to access services, as in some cloud deployment models, can significantly affect exposure to various threats and vulnerabilities. The potential service access configurations, i.e., via the use of local versus wide-area networks, are depicted in Figure 13. Scenarios (*a*) and (*b*) in Figure 13 represent current typical (non-cloud based) deployment models, whereby a service is hosted on an enterprise local-area network (LAN), and is either (*a*) accessed by a user via the LAN or (*b*) remotely using a virtual private network (VPN), for example. In terms of exposure to the vulnerabilities and threats associated with broad network access, which are shown in Table 5, arguably the *private cloud* deployment model, when services are hosted on the user’s premises to a single tenant, is the same as not using the cloud at all. Scenarios (*c*) and (*d*) are indicative of cloud deployments, whereby services exist remotely and are accessed via a wide-area network – these scenarios map to the private (single tenancy using remote hosting), community, public and hybrid cloud deployment models – and are arguably more exposed to the items shown in Table 5. The means of access to such deployments and the various entities involved are depicted in Figure 14. It illustrates the dependencies introduced on network providers and their infrastructure as well as the increase in attack surface. Distinguishing an attacker from a legitimate user on the cloud is a very challenging task, especially as the attacks can be launched from the Internet (a public resource) but also from within the same data centre (BNA-5).

Broad network access makes the aforementioned cloud deployment models more susceptible to attacks that exploit the vulnerabilities that are inherent in the Internet protocols (BNA-1) [GWS11]. For example, there are well-known security issues associated with the Border Gateway Protocol version 4 (BGPv4), which is responsible for inter-domain routing in the Internet [BFMR10]. Furthermore, given the pervasive nature of using wireless access network technologies, the security issues associated with wireless networks need to be considered, such as those discussed by Sheldon *et al.* [SWYP12]. For an indication of how network services typically fail, we refer the reader to ENISA’s “annual incidents report,” which summarises the

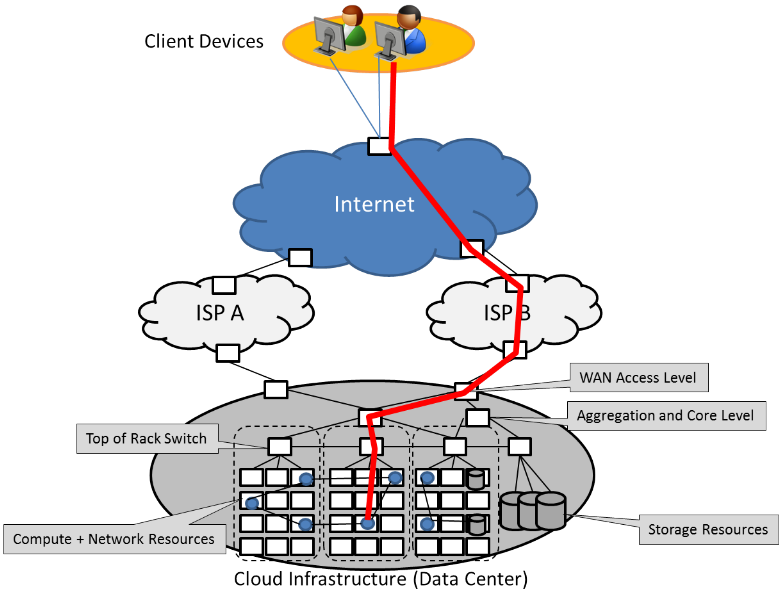
Client Devices
Internet
ISP A
ISP B
WAN Access Level
Aggregation and Core Level
Top of Rack Switch
Compute + Network Resources
Storage Resources
Cloud Infrastructure (Data Center)

Figure 14: Access to CI services on a public cloud infrastructure

Table 5: Broad network access-related vulnerabilities and threats

Broad Network Access				
ID	Description	Type	Obj	Refs
BNA-1	Inherent vulnerabilities in the Internet protocols	Vu	C-I-A	[BFMR10, SWYP12]
BNA-2	Denial of Service (DoS) attack against the cloud infrastructure via the public network infrastructure.	Th	A	[MR04, CXZB11]
BNA-3	Loss of wide-area network connectivity, e.g., due to misconfiguration, hardware failures, etc.	Vu	A	[cS13]
BNA-4	Man-in-the-middle attack on the user's public access network, e.g., a public WiFi infrastructure.	Th	C-I	[CCR09, WL03]
BNA-5	Denial of Service (DoS) attack against the CI service via the public network infrastructure or from within the data centre	Th	A	

significant incidents in electronic communications that are reported by Member States as part of implementing Article 13a [DKL13]. In a similar fashion to threats associated with on-demand self-service, broad network access introduces an increased potential for DDoS and MITM attacks (BNA-2 and -4), for example [MR04, CXZB11, CCR09]. In addition to the security issues associated with broad network access, cloud deployments that make use of wide-area networks, which are provided by third-parties, are arguably more susceptible to challenges such as misconfigurations and hardware failures, which can lead to a loss of connectivity (BNA-3). A taxonomy of such challenges has been developed by Çetinkaya and Sterbenz [cS13]. From an end-user perspective, these challenges can be difficult to diagnose [MSWA03], which could result in longer periods of service downtime.

6.1.3 Resource Pooling (RP)

Resource pooling relates to the ability of cloud infrastructure providers to provision their resources to multiple tenants, with physical and virtual resources being allocated based on consumer demand [MG11]. This capability is depicted in Figure 15. An implication of this is that consumers may not be completely aware of the precise location of services, as a number of data centres could be used to implement resource pooling by a provider. Also, it may be the case that a consumer will have limited knowledge of the other consumers they have been pooled with by a provider – some of these consumers may engage in nefarious activities, resulting in both technical and organisational threats. For example, a denial of resources attack could be instigated by a miscreant tenant via so-called API key misuse in IaaS and PaaS offerings (RP-1) [PMS13, Lem12]. A more subtle threat to data confidentiality may occur if suitable precautions are not taken to sanitise disk and volatile memory after a service (virtual machine) is terminated or migrated. Reconstructible remnants of data could be recovered by a malicious tenant that is subsequently allocated these memory resources (RP-3) [GWS11].

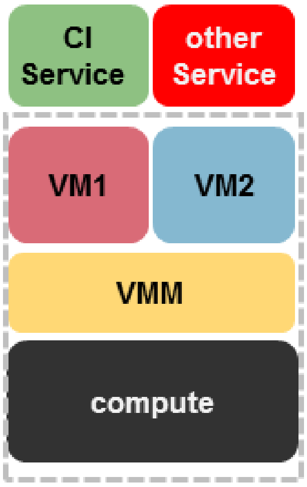

Figure 15: CI services sharing the compute node with other services and consumers

In many cases, services are constructed using a number of interacting components (or subservices). In the aforementioned box-model form of service deployment (see Figure 11), these components usually execute within the same trust domain, e.g., on the same physical device. However, in the cloud, these components may have to interact across different trust boundaries, e.g., across a data centre network, thus introducing the need for suitable security protection measures between interacting components. (This aspect of the transformation is depicted in Figure 16.) For example, this could include using IPSec between components or introducing firewalls – if these are not sufficiently provisioned, interactions between components represent a significant vulnerability (RP-8). Ways to identify trust boundaries, component interactions and the relevant threats can be discovered using structured analysis techniques, such as the Microsoft STRIDE method. An example of how this can be achieved is described by Klöti *et al.* in the context of a security analysis of OpenFlow networks [KKS13].

Another potential threat relates to the over-subscription of resources by the cloud infrastructure provider; these can include compute (RP-4), network (RP-5) and storage (RP-6) resources [BWT12]. This could occur because providers may wish to maximise the financial benefit from resource pooling, therefore running their infrastructure close to capacity, alongside unusual service demands.

Table 6: Resource pooling-related vulnerabilities and threats

Resource Pooling				
ID	Description	Type	Obj	Refs
RP-1	Denial of resources via miscreant API usage from a tenant	Th	A	[PMS13, Lem12]
RP-2	Jurisdictional collateral damage, e.g., shutdown request because of miscreant use from a malicious tenant	Th	A	[MS10, Vek13]
RP-3	Data recovery vulnerabilities, e.g., unauthorised access to data in memory or on disk from previous users	Vu	C-I	[GWS11]
RP-4 to 6	Over-subscription of the **compute** (4), **storage** (5), and **network** (6) resources by the cloud infrastructure provider	Th	A	[BWT12]
RP-7	Collateral damage from a tenant being attacked when insufficient tenant separation is employed	Th	A	[SKGK10, TJA10]
RP-8	Insufficient protection between virtual service (components) that must interact across "untrusted" regions of a data centre	Vu	C-I-A	[KKS13, H^+06]

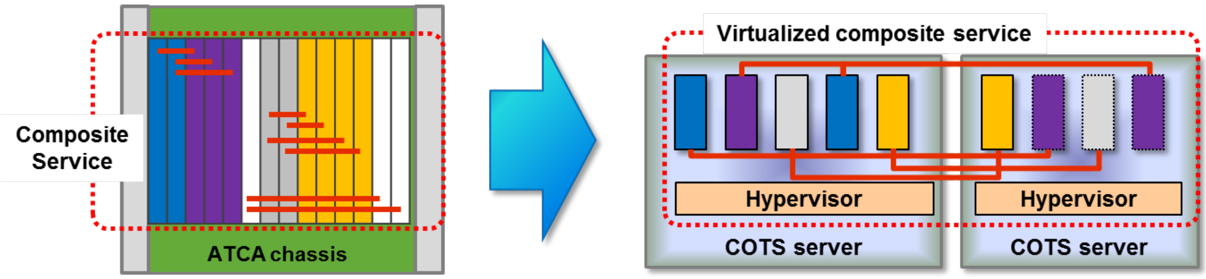

Figure 16: Exposure of formerly private interfaces to the outside

Furthermore, resource pooling introduces potential technical and organisational threats that are associated with so-called "collateral damage". From a technical perspective, this could include collateral damage from a DDoS attack targeted at (or originating from) a tenant that shares the same infrastructure – common resources, such as network access bandwidth, could be affected [SKGK10, TJA10] (RP-7). On an organisational level, Molnar and Schechter [MS10] highlight the potential for collateral damage occurring because of a "cease and desist" request from a law enforcement agency, resulting from miscreant behaviour by one of the tenants, to a cloud infrastructure provider (RP-2) [Vek13].

Of course, these vulnerabilities and threats are relevant only to the community, public and hybrid cloud deployment models, i.e., private cloud offerings, to a great extent, are not affected.

6.1.4 Rapid Elasticity (RE)

One of the benefits of cloud computing is the ability to add resources on demand when the load of the running service increases [MG11]. This ability is called *elasticity* and can be realised in two ways. Up-scaling is the dynamic addition of virtual resources by the VMM, e.g., providing additional CPU or memory to the VM. If the guest operating system (OS) within the VM supports scaling-up, these resources are immediately available to the service. If no more resources are available on the compute node or scaling-up is not supported by the guest OS, scaling-out is an alternative. Out-scaling means that additional resources are provided by a new VM. A service needs built-in functionality to use such distributed resources, e.g., load balancing, session scheduling, etc. The new VM can either be provided on the same compute node (see Figure 17(a) or on a remote compute node (see Figure 17(b).

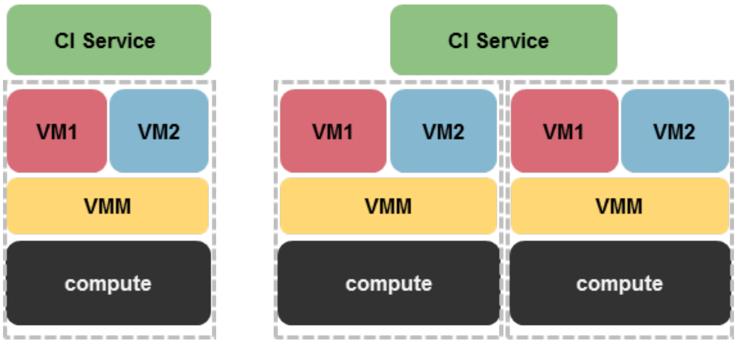

(a) Local scaling (b) Cloud scaling

Figure 17: Scaling out of services on the cloud

Whilst there are numerous benefits associated with elasticity, such as the ability to respond to transient peaks in user demand (e.g., a flash crowd), it introduces a number of potential vulnerabilities and threats. If it is assumed that resources will be available for scaling-up and -out, and if they prove not to be available, the availability of a service could be compromised. This is the case for compute (RE-1 and -4), network (RE-2 and -5) and storage (RE-3 and -6) resources. As mentioned earlier, scaling-out involves instantiating new virtual machines, on different hardware, in order to provide additional resources to a service (Figure 17(b)). This arrangement can introduce some problems if service deployment constraints are neglected in life-cycle management: certain Quality of Service (QoS) requirements may need to be fulfilled in the data centre network between components (which are executing on distinct hardware in a data centre), e.g., related to network delay. If these are not fulfilled, for example, because of poor virtual machine placement strategies, a service may fail to function correctly (RE-7). Similarly, virtual machines may become disconnected as a consequence of scaling-out, if they are placed behind a firewall that prohibits the required connectivity, for example (RE-8). Moreover, scaling-out of complex service functions requires more sophisticated scaling functions compared to state-of-the-art; for example, a firewall protecting a service function can be provisioned on a separate VM. If the service function is scaled-out but not the firewall, the service function may become vulnerable to attacks. Alternatively, if the scale-out does not respect the service function chain dependencies, the resulting deployment might deliver a degraded or incorrect service (RE-9).

Virtual service migration enables virtual machines to be moved between different underlying physical infrastructures. There are a number of approaches to virtual service migration, both within and across geographically (and topologically) distinct data centres, known as local and wide-area migration, respectively [FFCdM11, CFH$^+$05]. Migration has a number of benefits, including improved tolerance to hardware faults, natural phenomena such as floods, power outages, and performance benefits brought about by migrating service topologically closer to their clients [FFCdM11].

Despite these benefits, a number of technical and legal issues arise. In many countries and at a European level there are regulations that dictate the jurisdictional region that certain data,

Table 7: Rapid elasticity-related vulnerabilities and threats

Rapid Elasticity				
ID	Description	Type	Obj	Refs
RE-1 to 3	Insufficient underlying (1) compute, (2) network and (3) storage resources for **scaling-up**	Vu	A	[BWT12]
RE-4 to 6	Insufficient underlying (1) compute, (2) network and (3) storage resources for **scaling-out**	Vu	A	[BWT12]
RE-7	Scaling-out leads to performance issues because virtualised services that require certain network Quality of Service (QoS) requirements are not fulfilled	Th	A	
RE-8	Scaling-out leads to the disconnection of associated virtual machines	Th	A	
RE-9	Scaling-out violates service function chain dependencies and leads to unprotected instances or malfunction of service	Th	I-A	
RE-10	Failure of a migration process within a cloud data centre infrastructure (local-area migration)	Th	A	
RE-11	Failure of a migration process across different cloud data centre infrastructures (wide-area migration)	Th	A	
RE-12	Virtual service / data migration outside of accepted jurisdictional boundaries, e.g., critical infrastructure or eGovernment data leaving the EU.	Th	C	[BP13]
RE-13	Redundant service components are migrated to the same compute node; thus, they are subject to simultaneous failure.	Th	A	
RE-14	Information loss during VM migration.	Th	A	

e.g., related to the personal data of citizens, is allowed to reside. Unless explicitly requested, this may be challenging to track – these regulations could be breached when multi-national cloud providers are used (RE-12) [BP13].

From a technical perspective, virtual service migration may fail for a number of reasons, e.g., because of the issues highlighted in Section 6.1.2 relating to broad network access. This is possible for both local (RE-10) and wide-area (RE-11) migration of virtual services. Additionally, data may be lost during virtual service migration, which may not result in a service failure (RE-14). As a way of improving the fault-tolerance of a service and data, they may be provisioned in a redundant manner in distinct failure domains, e.g., on different racks in a data centre [PJGLSAH11]. However, if not configured correctly, redundant components may be located on the same compute node, thus resulting in a service failure if that node fails (RE-13).

6.1.5 Measured Service (MS)

The basis for the automatic control and optimisation that can occur in cloud offerings is a metering infrastructure, which measures usage at suitable levels of abstraction based on the type of service being provided [MG11]. This could, for example, include measuring storage, compute, network and the number of active user accounts. This metering infrastructure can underpin the billing of usage by cloud users, making it a potential target for attackers. For example, vulnerabilities in the metering infrastructure could be exploited, resulting in incorrect bills being generated for legitimate use (MS-1) [GWS11] – a challenge here is identifying that this has occurred and contesting bills as a cloud user. Conversely, cloud users could aim to subvert the metering infrastructure, in order to obtain reduced bills with respect to their actual consumption. A DoS attack could lead to an excessive use of metered cloud resources, which could be billed for, thus resulting in an *Economic* Denial of Service (E-DoS) [Dek12]. Furthermore, as there is likely to be an interface between the metering and billing infrastructures, attacks could occur that lead to a loss of confidentiality of account information (MS-2).

Table 8: Measures service-related vulnerabilities and threats

Measure Service				
ID	Description	Type	Obj	Refs
MS-1	Vulnerabilities in the measurement infrastructure that is used for billing	Vu	I	
MS-2	Attacks to the metering infrastructure that holds confidential account information, e.g., for payment	Th	C-I	
MS-3	An attack that leads to an excessive use of resources that are billed for, resulting in an Economic Denial of Service (E-DoS)	Th	A	[Dek12]

6.1.6 Virtualisation (VN)

Virtualisation is arguably one of the key technologies that enables cloud-based service provisioning [MG11]. For instance, it is the technology that supports rapid elasticity. Whilst there are numerous benefits associated with virtualisation, it adds a number of new vulnerabilities and potential threats. The VMM, which supports activities such as scaling-out and service migration, is central to the management of virtual machines that execute in the cloud. Potential software vulnerabilities in the VMM could be exploited by an attacker, with wide-ranging security implications (VN-1), depending on the nature of the exploit. For example, *virtual machine escape* threats have been demonstrated, whereby processes can gain unauthorized access to state outside their virtual machine (VN-4) [Hig09]. Similarly, side-channel attacks on VMs hosted on the same compute node have been demonstrated (VN-5) [ZJRR12]. Furthermore, failures of the VMM for non-malicious reasons, such as software bugs (VN-2) or misconfiguration by the cloud operator (VN-3), can significantly affect a virtual service executing on it.

In non-virtualised environments hardware watchdogs, for example, are used to detect component failure and to restart the appliance. That is not feasible in cloud environments as the watchdogs cannot restart and configure a VM image. Moreover, in some cases, failures of the underlying infrastructure can become transparent to the virtualisation environment, thus rendering failure detection and service remediation methods ineffective (VN-6).

Table 9: Virtualisation-related vulnerabilities and threats

Virtualisation				
ID	Description	Type	Obj	Refs
VN-1	Software vulnerabilities in the Virtual Machine Manager (VMM)	Vu	C-I-A	
VN-2	Failure of the Virtual Machine Manager (VMM) because of software bugs	Th	A	
VN-3	Failure of the Virtual Machine Manager (VMM) because of misconfiguration by the cloud infrastructure provider	Th	A	
VN-4	Processes escaping from a virtual machine – virtual machine escape	Th	C-I	[Hig09]
VN-5	Cross virtual machine side-channel attack	Th	A	[ZJRR12]
VN-6	Failures of the physical infrastructure are transparent to the virtual environment	Th	A	

6.1.7 Organisational Issues (OI)

By deploying services in the cloud, a number of organisational threats and vulnerabilities can emerge, as discussed by Molnar and Schechter [MS10].

In general, enterprise organisational boundaries are becoming increasingly blurred, e.g., with the use of subcontractors for provisioning services, which makes understanding the *insider threat* challenging [FvCvEW10]. Insider threats are those that are realised by disgruntled or negligent employees, for example. The Carnegie Mellon CERT maintains a useful knowledge base regarding the insider threat[7]. Using the cloud makes determining the risk of threats from insiders challenging, as one must attempt to consider potential adversaries associated with a number of enterprises that are used to provide a service, not all of which are transparent to the cloud user (OI-1). Closely related to the issue of the insider threat is the poor implementation of information security processes by the cloud provider (OI-2) [MS10]. Industry-accepted certification, such as ISO 27000, can be used to indicate whether appropriate information security management processes are being implemented by an organisation.

When a security incident occurs, it is necessary to have incident-response management strategies in place that can be used to localise a problem, determine the impact associated with an incident, and, if necessary, inform end-users of data breaches. For various reasons, incident-response management in non-cloud settings is challenging, especially when attacks are targeted and use advanced malware [Ree10]. However, this problem is further compounded in a cloud context, wherein digital forensics capabilities may be restricted and activities must be coordinated with a third-party organisation, such as the cloud infrastructure provider [THGL11]. Therefore, suitable incident-response management plans must be in place – if not, an incident may be further compounded and the potential impact made worse (OI-3).

Related to incident-response management are issues associated with poor Service Level Agreement (SLA) specification for the cloud. In addition to clearly defining performance requirements in an SLA, there must be a clear specification of responsibilities in relation to security – without this vulnerabilities could occur (OI-4) [ANM11].

[7] http://www.cert.org/insider_threat/

Table 10: Organisational issues-related vulnerabilities and threats

Organisational Issues				
ID	Description	Type	Obj	Refs
OI-1	Malicious insiders in the cloud infrastructure provider	Th	C-I-A	[FvCvEW10] [ANM11]
OI-2	Poor information security processes operated by the cloud infrastructure provider	Vu	C-I-A	
OI-3	Inadequate incident-response management processes by the cloud infrastructure provider	Vu	C-I-A	[Ree10]
OI-4	Issues emerging because of poor SLA specification	Vu	C-I-A	[ANM11]
OI-5	Contractual issues that emerge because of bankruptcy and potential switching costs	Vu	A	[MS10] [JKKD12]
OI-6	Failure of a sub-contractor, which is used by the primary “obligor,” i.e., a cloud infrastructure provider	Vu	A	[BP13]
OI-7	Vulnerabilities emerging from a lack of control of software versions and APIs	Vu	C-I-A	
OI-8	Reduction of in-house expertise caused by outsourcing services, resulting in a lack of organisation resilience when challenges occur, such as attacks.	Vu	A	
OI-9	Misuse of an organisation’s data, as specified in the terms of use, e.g., for advertising or resale.	Th	C	

A significant threat in the cloud can occur, which does not directly relate to cyber-security, when a provider organisation becomes bankrupt and switching costs are high, thus rendering migration to a new infrastructure provider problematic (OI-5) [MS10]. Consequently, a critical infrastructure service could experience an extended period of unavailability. Measures to indicate the probability of this relate to the longevity of the organisation used to provide a service, their reported financial status, and the use of open standards to support switching to a different provider, if necessary. These organisational measures are discussed by the OPTIMIS project, in the risk assessment deliverable [JKKD12].

Cloud services can be provided by a number of organisations, in addition to those that have a direct user and provider relationship. For example, a cloud infrastructure provider may subcontract to other providers, e.g., to implement scaling-out. The use of multiple providers is likely to become more prevalent with so-called “cloud of cloud” offerings, whereby cloud services, such as data storage, are offered using a number of providers, which are administered by brokers [Blo11]. A challenging threat to address can occur when an organisation that has been sub-contracted fails to provide a service (OI-6). SECCRIT deliverable D2.2 discusses legal aspects associated with this threat [BP13].

When using IaaS and PaaS cloud offerings, a threat could occur when providers change the APIs they use, which a critical infrastructure service provider’s software depends on (OI-7). This could lead to a service becoming unavailable whilst updates are made to ensure compatibility with new APIs, or more subtly lead to vulnerabilities because of inappropriate interpretation of the functionality of a new API.

Migrating services to the cloud has the benefit of reducing the need to maintain in-house expertise regarding the provisioning of certain aspects of a service. Whilst this can help to reduce an organisation's operational costs, it may lead to a lack of organisational resilience when incidents occur, or the inability to migrate services in-house if problems occur with a cloud offering (OI-8). This is a longer-term threat that may emerge from using the cloud.

In many cases, consumer-oriented cloud offerings aim to generate revenue from the data they garner from their users, e.g., for enabling targeted advertisements, product endorsements, or direct reselling of (marketing/meta) data. A significant organisational threat relates to the misuse of an organisation's data, as specified in the terms of use, for these purposes (OI-9).

6.1.8 Threats and Vulnerabilities Prevalent in Cloud Offerings (CO)

There are threats and vulnerabilities that are prevalent in contemporary cloud offerings, which are noteworthy. Many cloud offerings make use of web-based technologies to implement SaaS solutions and for their management interface. Consequently, they are potentially vulnerable to threats that emerge from the use of web-based technologies, such as SQL injection attacks [Aug09] and cross-site scripting [ST12] (CO-1) [GWS11, SK11]. Furthermore, it is understood that there have been vulnerabilities in the APIs that are used to build cloud services – this threat appears in the Cloud Security Alliance's "Notorious Nine" (CO-2) [CSA13].

Table 11: Contemporary cloud offering-related vulnerabilities and threats

Security Problems Prevalent in Cloud Offering				
ID	Description	Type	Obj	Refs
CO-1	Threats that emerge from the use of Web-based technologies, such as SQL injection attacks and cross-site scripting	Th	C-I-A	[GWS11, SK11]
CO-2	Insecure interfaces and APIs that are used to build cloud services	Vu	C-I-A	[CSA13]

In general, we recommend the reader to stay informed about these threats via the Cloud Security Alliance's threat catalogue (or other similar advisories), which is updated annually and is based on expert opinion in the field of cloud computing security. Furthermore, it contains pointers to relevant information, including references to protection measures that can be applied to the threats.

6.1.9 Physical Infrastructure (PI)

We have included a number of threats that are related to the underlying physical data centre infrastructure. We have included these threats for two primary reasons: *(i)* in contrast to traditional critical infrastructure services, which tend to be implemented on hardware with fault-tolerant mechanisms built-in (e.g., redundant power supplies), cloud offerings are normally provisioned on unreliable hardware (with virtualisation masking hardware failures); and *(ii)* decisions regarding how to manage threats to the underlying infrastructure are typically beyond the control of the cloud user, and should therefore be assessed with respect to the (additional) risk they introduce.

Failure of the physical compute (PI-2), storage (PI-3) and network (PI-4) facilities due to hardware faults or misconfiguration represents a common threat in large data centres that support cloud services. In many cases, hardware failures are not perceptible from the service level because of virtualisation and the inherent redundancy in data centres. However, a number of high-profile cloud outages have resulted from misconfiguration, largely in the network infrastructure [Ver13, Nei12].

Table 12: Physical infrastructure-related vulnerabilities and threats

Physical Infrastructure				
ID	Description	Type	Obj	Refs
PI-1	Loss of power to the cloud data centre, e.g., caused by a regional outage or surge	Th	A	[Tal13]
PI-2 to 4	Failure of physical *compute* (2), *storage* (3) and *network* (4) facilities in the cloud data centre	Th	A	[Ver13, Nei12]
PI-5	Unauthorised physical access to the cloud data centre infrastructure	Th	C-I-A	[Bar12]
PI-6	A (large-scale) catastrophe, such as a fire or earthquake, that affects the cloud data centre	Th	A	[Ama13, Alu12]

Failures caused by power outages to the cloud data centre can have a significant impact on cloud-based services (PI-1) [Tal13]. In these cases, backup power facilities should be in place to maintain the availability of services (or services such as those provided by Amazon, which enable regional diversity should be employed [Ama13]). Other natural phenomena, such as floods, fires and earthquakes, may cause a loss of availability of a cloud data centre (PI-6). Conversely, if applied appropriately, the use of cloud computing could result in improved resilience to such events [Alu12]. Care should be taken that data centres are located in areas that are not susceptible to such problems, e.g., they are not constructed on floodplains. Large cloud providers, such as Amazon, provide services that support geographical diversity of data centres [Ama13].

Another important threat relates to physical infrastructure security (PI-5) [Bar12], i.e., ensuring that unauthorised physical access to the data centre infrastructure does not occur. In general, this issue relates to the information security management processes that are in place at a cloud provider, e.g., related to the implementation of the ISO 27000 standard.

6.1.10 Security Control Implementation (SCI)

Cloud computing can make the implementation of existing security controls challenging to implement, which introduces a number of vulnerabilities [GWS11]. Cloud services often run across multiple sites, which can make key management more complicated; furthermore, studies have shown that cloud services typically require a large number of keys, exacerbating this problem (SCI-1) [ENI12].

Monitoring network behaviour is a key security control, including the detection of anomalies which could indicate attack behaviour or misconfiguration, for example. The aforementioned rapid elasticity property of clouds has the potential to reduce the performance of contemporary anomaly detection techniques, which examine network traffic, thus resulting in unacceptable (and insecure) false positive and negative rates (SCI-2). Initial studies have shown this to be the case [Ada13]. These issues will be investigated further in the SECCRIT project in work package four.

An important way of understanding the security vulnerabilities associated with services is to run network-based vulnerability analysis tools, such as OpenVAS[8]. For a number of reasons, including the unwillingness of tenants, it can be challenging to execute these tools in a cloud environment – in many cases, these tools have a behaviour “profile” that is similar to actual attack behaviour (SCI-3). Consequently, the vulnerabilities associated with a service, whilst it operates in the cloud, can be difficult to determine, thus resulting in potentially exploitable vulnerabilities.

Table 13: Control implementation-related vulnerabilities and threats

Security Control Implementation Challenges				
ID	Description	Type	Obj	Refs
SCI-1	Key management issues, brought about due to multiple geographical locations of cloud services; many keys can be required	Vu	C-I-A	[ENI12]
SCI-2	Network-based anomaly detection is not readily implementable	Vu	C-I-A	[Ada13]
SCI-3	Inability to run network-based vulnerability scanning and penetration testing tools	Vu	C-I-A	[MS10]
SCI-5	Widely used control mechanisms not available to virtualised environments	Vu	A	

A particular challenge of transferring services to a virtualised environment is the inability to reuse widely used control mechanisms. For example, in telecommunications, to distinguish the failure modes “dead peer” and “communication failure” a serial line between two Advanced Telecommunications Computing Architecture (ATCA) boxes is used to have independent ways to test the failure mode. In a virtual environment, a serial line connection to a peer instance would be mapped onto a virtual link; this virtual link would use the same physical infrastructure as the communication link between the two instances, thus making it impossible to determine the correct failure mode.

6.2 Summary

In this section, we have presented a cloud-specific vulnerability and threat catalogue that can be applied when an organisation implements a risk assessment in this context. In order to ensure the items in the catalogue relate specifically to cloud issues, they are organised into a number of categories that are associated with NIST’s essential cloud computing characteristics and associated issues, such as the properties of key technologies. The threats and vulnerabilities are presented at a level of abstraction that enables a risk analyst to consider them in a deployment independent manner, and then apply them to a specific implementation. In other words, we have avoided presenting issues associated with specific products, for example. Consequently, in many cases, the catalogue should be augmented (or instantiated) with issues that relate to a specific deployment with items that are determined from knowledge bases such as the NIST national vulnerability database[9], security advisories, and vulnerability assessment

[8] `http://openvas.org/`

[9] http://nvd.nist.gov/

tools. In Section 7, we indicate how the catalogue can be used by an organisation that supports high-assurance ICT services to determine the risks of adopting the cloud.

7 Determining the Risk of Cloud Adoption

A significant decision that organisations which operate high assurance ICT services must face is whether to adopt cloud computing and, more specifically, which configuration of cloud offering they should use, e.g., a private, public or hybrid cloud. In this section, we outline a risk-based approach that organisations can use to support this decision making process. The approach outlined here can be used alongside others, such as the one proposed by the OPTIMIS project [JKKD12] to support the analysis of different cloud offerings. In addition, there are commercial services that aim to support organisations with this decision making process, which include some security factors, such as whether the cloud provider has security certification.

We assume organisations that are considering migrating some of their data and services to the cloud already conduct an information security risk assessment process, e.g., in order to comply with the ISO 27001 information security management system standard. This is a reasonable assumption, especially for organisations that operate high assurance ICT services. Furthermore, this assumption will become even more valid in the future if the proposed EU Network and Information Security (NIS) Directive [Eur13] comes into force, which states:

(22) *"...A culture of risk management, involving risk assessment [...] should be promoted and developed through appropriate regulatory requirements and voluntary industry practices..."*

(24)*"...Those obligations should also be extended to public administrations, and operators of critical infrastructure which rely heavily on information and communications technology [...] such as electricity and gas, transport, credit institutions, stock exchange and health..."*

The approach we propose builds on the outcomes of the existing risk assessment conducted by an organisation, and augments it by considering the risk scenarios, i.e., threats and vulnerabilities, that relate to the use of the cloud. For the purpose of exposition, we show how this can be achieved using the risk assessment process that is supported by the open-source Verinice ISMS tool[10], in the context of one of the SECCRIT project scenarios.

7.1 Risk Assessment Scenario

The scenario we will use to describe our approach to determining risks from cloud usage is based on a video surveillance system. The surveillance system includes a number of ICT assets that could be migrated to the cloud, including services and data, which can be summarised, as follows:

- *Live video data*: This is live video footage taken from the surveillance cameras at a deployment location, which is being analysed in order to detect anomalies that could indicate an incident is occurring.
- *Archive video data*: This is archive footage from the surveillance cameras that needs to be stored, either because it was asked for by a client organisation or is needed to fulfil a regulatory obligation.

[10] `http://www.verinice.org`

- *Anomaly detection services*: These are the processes, realised as applications, which are used to detect anomalous behaviour using the live video data. For example, such algorithms can be used to detect people moving quickly against the flow of a crowd, or detect activity at unusual times of the day.
- *Data archiving services*: These are the processes that are used to prepare the live video footage for archiving. For example, this could include compressing and encrypting the data for it to be stored on disk. Also, this processing might include the generation of meta-data.
- *Security operator interface*: This is the interface that security operators receive and inspect the security alerts that are generated by the anomaly detection services. This could be implemented using a 'dumb terminal', which connects to a server in the cloud, for example.

These ICT assets, including an overview of the architecture and the different stakeholders associated with a potential cloud deployment, are depicted in Figure 18. Currently, the ICT assets are deployed at the customer premises (the area shown under TenSys region in Figure 18) and are solely operated by TenSys – the company responsible for operating the surveillance system. Migrating some of these services and data to the cloud has a number of potential benefits, including reducing the amount of on-site infrastructure, the ability to scale resources based on demand, e.g., in response to a public safety incident, and being able to rapidly accommodate new installations. However, it is unclear what the risks of adopting cloud computing are, and what form of offering should be adopted.

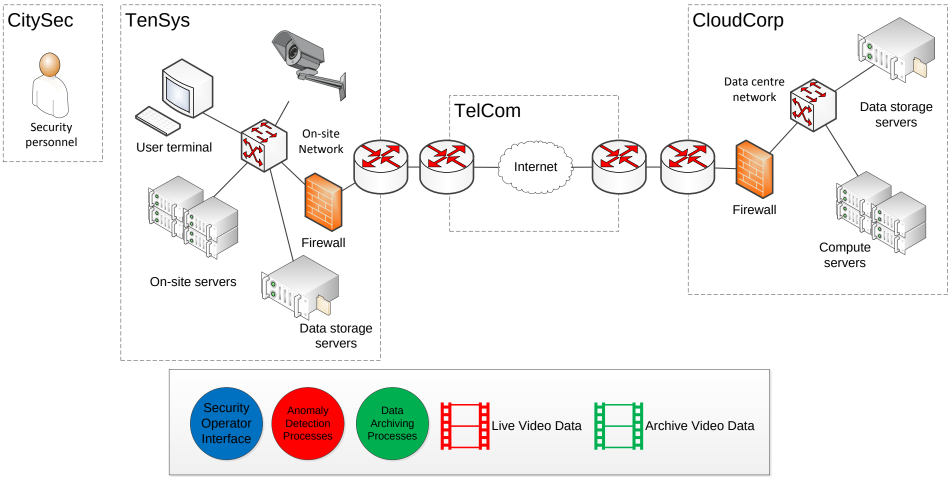

Figure 18: The ICT assets, infrastructure and stakeholders that will be considered in the example cloud adoption scenario

We assume the organisation that is responsible for operating the video surveillance system (TenSys) has conducted an information security risk assessment in relation to these ICT assets. In the following section, we introduce one approach they could use to achieve this, using the Verinice ISMS tool, and how it can be extended to determine cloud-related risks.

7.2 Cloud Adoption Risk Assessment Process

The process that an organisation can follow to determine the risks associated with cloud adoption is summarised in Figure 19. It consists of two parts: *(i)* the standard Verinice-supported information security risk assessment process[11]; and *(ii)* an extension that augments the outcomes of this assessment to consider a potential cloud deployment scenario. It is worth noting that the basic process supported by Verinice is similar in nature to other approaches, whose focal point are the assets associated with an organisation; such approaches are outlined in Section 2. Arguably, therefore, the extension we propose can be applied to any risk assessment process that takes an asset-driven approach. We assume organisations that are considering migrating some of their ICT assets to the cloud have available to them a set of offerings, e.g., from cloud infrastructure providers, which they are considering – these are modelled in Stage 8 of the process. In what follows, we outline how the process defined in Figure 19 can be applied to the video surveillance system scenario, depicted in Figure 18.

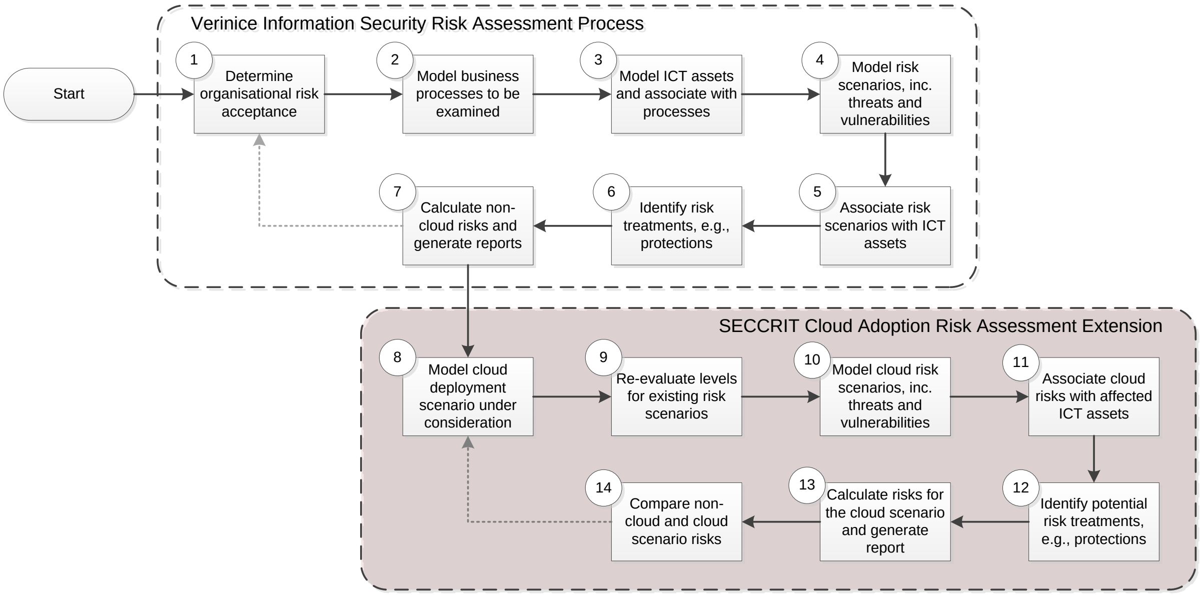
Verinice Information Security Risk Assessment Process
Start
1 Determine organisational risk acceptance
2 Model business processes to be examined
3 Model ICT assets and associate with processes
4 Model risk scenarios, inc. threats and vulnerabilities
5 Associate risk scenarios with ICT assets
6 Identify risk treatments, e.g., protections
7 Calculate non-cloud risks and generate reports
SECCRIT Cloud Adoption Risk Assessment Extension
8 Model cloud deployment scenario under consideration
9 Re-evaluate levels for existing risk scenarios
10 Model cloud risk scenarios, inc. threats and vulnerabilities
11 Associate cloud risks with affected ICT assets
12 Identify potential risk treatments, e.g., protections
13 Calculate risks for the cloud scenario and generate report
14 Compare non-cloud and cloud scenario risks

Figure 19: Overview of the Verinice information security risk assessment process and the SECCRIT cloud adoption risk assessment extension

7.2.1 Verinice Information Security Risk Assessment

The first stage of the risk assessment process outlined in Figure 19 is to define an organisation's overall risk acceptance with respect to the *confidentiality*, *integrity* and *availability* security objectives (see Section 4 for a discussion on these). These values are used by Verinice to determine the overall risk levels for these factors when they are calculated at a later stage. For TenSys – the operator of a video surveillance system – arguably, the confidentiality of data is of paramount importance, closely followed by the integrity and availability of data and associated services. Therefore, confidentiality is given a relatively low acceptance value, as can be seen in Figure 20.

Having entered the organisation's overall risk acceptance levels, the next step in the Verinice risk assessment is to model the processes that relate to an organisation (or system) under consideration. These will have ICT assets associated with them in a next step. In the context of the video surveillance case study, the following processes can be identified:

[11] In what follows, we make use of the Verinice software version 1.6.3

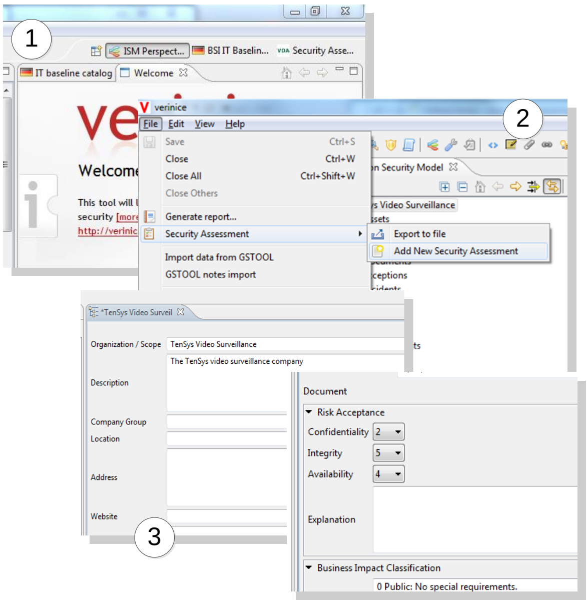

Organization / Scope: TenSys Video Surveillance

Description: The TenSys video surveillance company

Company Group

Location

Address

Website

Confidentiality: 2

Integrity: 5

Availability: 4

Business Impact Classification: 0 Public: No special requirements.

Figure 20: Creating a new security assessment, entering organisation details and risk acceptance levels

To create a new security assessment in Verinice, from the `File` menu, under the `Security Assessment` option, select `Add New Security Assessment` (2). This will create a new template for an organisation's security assessment. From the *ISM Perspective* (1), an organisation's details can be entered, including the risk acceptance values for confidentiality, integrity and availability (3).

- *Detecting anomalous behaviour*: TenSys have developed a number of algorithms and processes that can be used to detect anomalous behaviour in video surveillance footage, e.g., persons moving in the opposite direction to the crowd – the premise being such behaviour could be indicative of malicious activity.
- *Notifying security personnel of an incident*: having detected anomalous behaviour that could be indicative of nefarious behaviour, this process relates to notifying security personnel (under the employment of CitySec) that their attention is required.
- *Archiving of video surveillance footage*: this is the process of taking video surveillance data, and preparing and storing it for long-term archiving. This is a process that could be requested by the client or needed to meet regulatory obligations.
- *Searching archival video surveillance footage*: this is the process that supports post-mortem searching of video surveillance data, in order to determine the cause of an incident. The process builds on the archived surveillance data.

The key objective with this stage, looking ahead to determining cloud-associated risks, is to model the organisational processes that could be, in part, supported by cloud computing. These may be closely coupled to the specific domain the organisation is operating, such as those described above, or processes such as invoicing or customer-relations management. A further step that is required is to determine the *business impact* that could occur if these processes are interrupted. The steps to achieve this in the Verinice tool are outlined in Figure 21.

The defined processes are supported by a number of assets, such as information, software, people and hardware, which are modelled next. In the video surveillance case study, these are the assets that are shown in Figure 18. The steps for creating these assets and associating them with organisational processes in the Verinice tool are shown in Figure 22.

The next stage in the process is to define the *risk scenarios* that could affect the assets that have been identified. These scenarios are associated with threats and vulnerabilities, which must be assessed with respect to their *frequency* and *severity*, respectively. Developing these

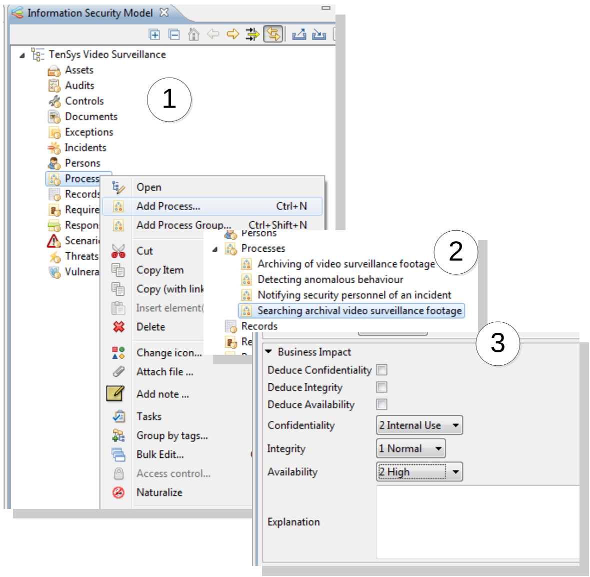

Deduce Confidentiality ☐
Deduce Integrity ☐
Deduce Availability ☐
Confidentiality: 2 Internal Use
Integrity: 1 Normal
Availability: 2 High

To create a new process, right click on the `Processes` folder, which is listed under the organisation details for the security assessment, and select `Add Process...` from the context menu (1). Fill-in the details for the new process that you have created, e.g., its name and so on. Repeat this process for each of the processes that have been identified (2). Finally, for each of the processes, define the business impact (in terms of confidentiality, integrity and availability) if these processes are interrupted (3).

Figure 21: Modelling processes and their impact on the business if they are interrupted

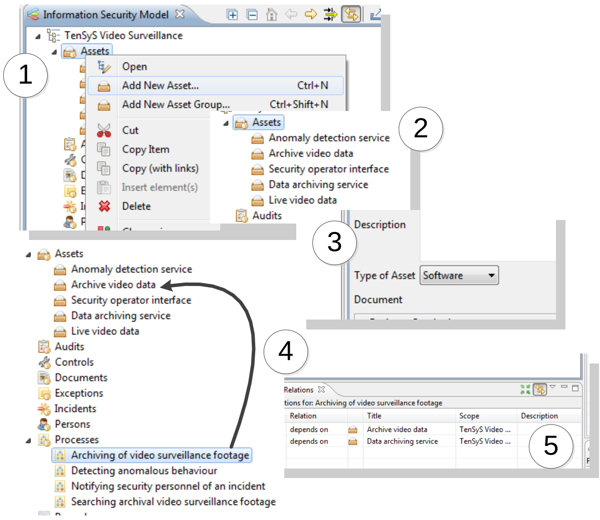
Information Security Model
TenSyS Video Surveillance
Assets
1
Open
Add New Asset... Ctrl+N
Add New Asset Group... Ctrl+Shift+N
Cut
Copy Item
Copy (with links)
Insert element(s)
Delete
2
Assets
Anomaly detection service
Archive video data
Security operator interface
Data archiving service
Live video data
Audits
3
Description
Document
Assets
Anomaly detection service
Archive video data
Security operator interface
Data archiving service
Live video data
Audits
Controls
Documents
Exceptions
Incidents
Persons
Processes
Archiving of video surveillance footage
Detecting anomalous behaviour
Notifying security personnel of an incident
Searching archival video surveillance footage
4
Relations
5

Type of Asset: Software

tions for: Archiving of video surveillance footage

Relation		Title	Scope	Description
depends on		Archive video data	TenSyS Video ...	
depends on		Data archiving service	TenSyS Video ...	

To create a new asset, right click on the `Assets` folder, which is listed under the organisation details for the security assessment, and select `Add New Asset....` Fill-in the details for the asset. Repeat this process for each of the assets that have been identified (2). For each of the assets that have been created, define the type of asset that it is, e.g., software, information, people, etc. (3). These assets support the organisational processes that were previously defined, and need to be associated with each other. This can be achieved by dragging a process onto the asset that supports it (4). A "depends on" relationship should then exist between the process and asset (5).

Figure 22: Modelling the assets associated with the organisation's processes, and associating them

risk scenarios, and determining the aforementioned values, is arguably one of the more challenging aspects of a risk assessment. To support this stage in the risk assessment process, the professional version of the Verinice tool includes a standard catalogue of risk scenarios, which can be acquired separately. (This should be augmented with risk scenarios that are specific to an organisation.) Furthermore, the German Federal Office for Information Security (BSI) IT Baseline Catalogue can be imported into the Verinice tool[12], which can support the development of risk scenarios. However, at the time of writing, the catalogue is only available in the German language. Furthermore, the results for vulnerability scanning tools, which have been applied to the target system, can be used to populate the vulnerability catalogue;

[12] The catalogue can be downloaded from the Verinice website at: `http://www.verinice.org/download/`

the open-source vulnerability scanning tool OpenVAS[13] supports the exporting of its results to Verinice, for example. The process for developing risk scenarios in the Verinice tool is outlined in Figure 23.

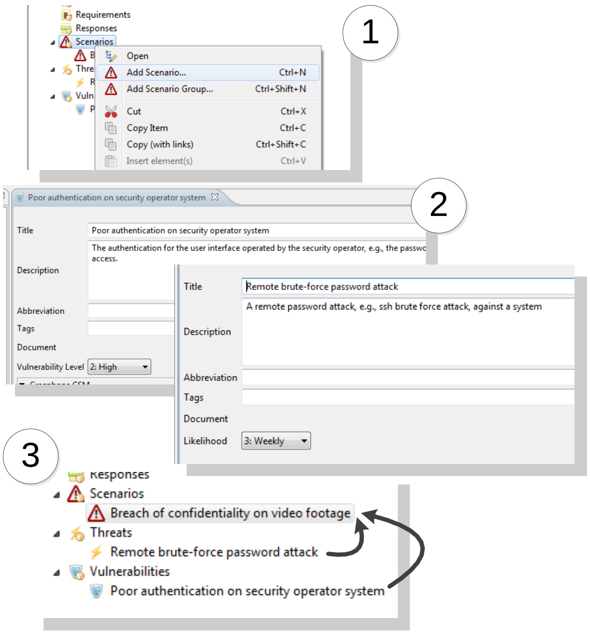
1
Requirements
Responses
Scenarios
Threats
Vulnerabilities
Open
Add Scenario... Ctrl+N
Add Scenario Group... Ctrl+Shift+N
Cut Ctrl+X
Copy Item Ctrl+C
Copy (with links) Ctrl+Shift+C
Insert element(s) Ctrl+V
2
Poor authentication on security operator system
3
Responses
Scenarios
Breach of confidentiality on video footage
Threats
Remote brute-force password attack
Vulnerabilities
Poor authentication on security operator system

Title: Poor authentication on security operator system
Description: The authentication for the user interface operated by the security operator, e.g., the passwc access.
Abbreviation:
Tags:
Document:
Vulnerability Level: 2: High

Title: Remote brute-force password attack
Description: A remote password attack, e.g., ssh brute force attack, against a system
Abbreviation:
Tags:
Document:
Likelihood: 3: Weekly

To create a new risk scenario, in a similar fashion to the previously modelled entities, right-click on the `Scenarios` folder, under the organisation details, and select `Add Scenario...` (1). A similar process can be followed for creating new threats and vulnerabilities. A description of the scenario, threats and vulnerabilities can be given. The threats should be evaluated with respect to their frequency and the vulnerabilities should have their severity rated (2). Finally, the threats and vulnerabilities can be associated with a scenario by dragging and dropping them in the Verinice tool onto the relevant scenarios (3). Note that threats and vulnerabilities may be associated with more than one scenario.

Figure 23: Modelling risk scenarios, including the threats and vulnerabilities

Having defined the risk scenarios, they must be associated with the relevant assets that have been previously defined. In the Verinice tool, this can be trivially achieved by dragging and dropping the scenarios onto the relevant assets, thus creating a relationship between them, which will be used when risk is calculated. This process is depicted in Figure 24; the resulting relationships between the assets, the scenarios and processes are also shown.

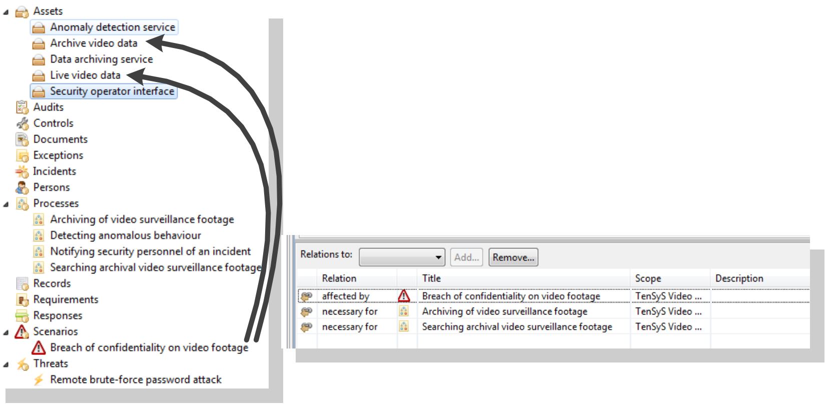
Assets
Anomaly detection service
Archive video data
Data archiving service
Live video data
Security operator interface
Audits
Controls
Documents
Exceptions
Incidents
Persons
Processes
Archiving of video surveillance footage
Detecting anomalous behaviour
Notifying security personnel of an incident
Searching archival video surveillance footage
Records
Requirements
Responses
Scenarios
Breach of confidentiality on video footage
Threats
Remote brute-force password attack
Relations to:
Add...
Remove...

Relation		Title	Scope	Description
affected by	⚠	Breach of confidentiality on video footage	TenSyS Video ...	
necessary for		Archiving of video surveillance footage	TenSyS Video ...	
necessary for		Searching archival video surveillance footage	TenSyS Video ...	

Figure 24: Associating the risk scenarios with the assets associated with the organisation

A number of different security controls can be put in place that affect the scenarios that have been previously defined. For example, they either reduce the probability that a scenario will occur or reduce its impact on confidentiality, integrity and availability. These controls can be both technical and organisational in nature. This stage of the risk assessment process is concerned

[13] The OpenVAS open-source vulnerability scanning tool: `http://www.openvas.org`

with identifying the security controls that have been implemented within an organisation, and associating them with and determining their affect on the previously defined risk scenarios. In a similar way to the standard risk scenarios that are available with the professional version of Verinice, there is a catalogue of controls based on the ISO 27002 standard available. Also, the BSI IT Baseline Catalogue includes a number of protection measures that can be applied to the security challenges that are defined within it. These standard security controls should be complemented with those that are defined in a specific organisational context. The steps to implement this part of the risk assessment in Verinice are outlined in Figure 25.

To create a new security control, right-click on the `Controls` folder under the organisational details, and select `Add Control...`; this will create a new security control, whose description can be filled-in (1). An essential task is to determine whether the control is implemented (2), and the *strength* of the control (3). The strength relates to its impact on the CIA security objectives and the probability of an associated scenario occurring. Each control can then either be associated with a risk scenario or an asset, by dragging and dropping it onto the related item (4). When this has been completed, a relationship is shown between the control and the scenario that it relates to (5).

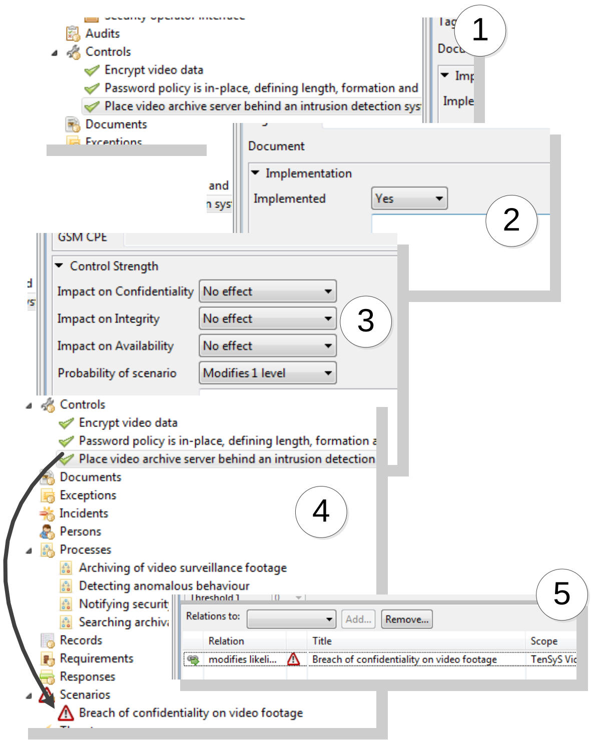

Audits
Controls
- Encrypt video data
- Password policy is in-place, defining length, formation and
- Place video archive server behind an intrusion detection sys

Documents

Document
Implementation
Implemented: Yes

Control Strength
Impact on Confidentiality: No effect
Impact on Integrity: No effect
Impact on Availability: No effect
Probability of scenario: Modifies 1 level

Controls
- Encrypt video data
- Password policy is in-place, defining length, formation a
- Place video archive server behind an intrusion detection

Documents
Exceptions
Incidents
Persons
Processes
- Archiving of video surveillance footage
- Detecting anomalous behaviour
- Notifying securit
- Searching archiv

Records
Requirements
Responses
Scenarios
- Breach of confidentiality on video footage

Relations to:

Relation		Title	Scope
modifies likeli...	⚠	Breach of confidentiality on video footage	TenSyS Vic

Figure 25: Defining and associating controls that affect scenarios and assets

Having defined the relevant security controls and how they augment the scenarios, the overall risk assessment can be carried out and a suitable report can be generated. In the Verinice tool, this is realised by initiating a risk analysis and subsequently generating a report, as outlined in Figure 26. There are a number of different reports that can be generated by the Verinice tool that shown different aspects of the risk posture of an organisation. For example, a report can be generated that shows the overall risk situation for an organisation, the assets that have a medium-to-high risk value associated with them, the high risk areas and the numbers associated with different security objectives, risk matrices that are similar to those proposed by the ISO 27005 standard, and detailed risk tables that enable closer inspection of the different risks.

As mentioned earlier, we assume that an organisation has carried out a similar risk assessment process to the one outlined here. This assessment forms the basis of our approach to determining the risks associated with cloud adoption, which we will discuss next.

7.2.2 SECCRIT Cloud Adoption Risk Assessment Extension

Building on the base risk assessment, the first step to follow (Step 8 in Figure 19) in order to determine the risks associated with cloud adoption is to model the cloud scenario(s) that

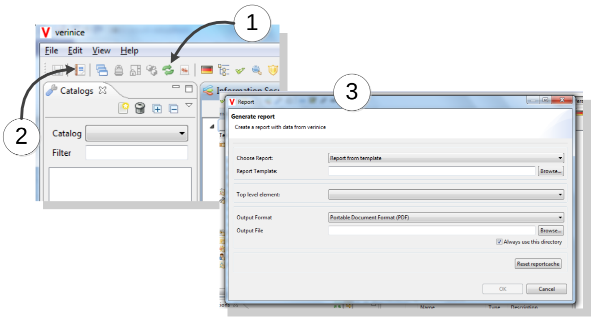

Generate report

Create a report with data from verinice

Choose Report: Report from template

Report Template: Browse...

Top level element:

Output Format: Portable Document Format (PDF)

Output File: Browse...

☑ Always use this directory

Reset reportcache

To run a risk analysis in the Verinice tool, select the `Run risk analysis` button from the menu bar (1). Subsequently, a risk assessment report can be generated, using the `Generate report...` button (2). There are a number of different report templates that can be used, in a number of different file formats (3).

Figure 26: The steps for executing a risk analysis and generating a risk assessment report

are under consideration. As mentioned earlier, the assumption is that an organisation has investigated (or has been offered) a set of cloud offerings they are considering, and that the organisation needs to understand the risks of adopting these offerings.

Model the cloud deployment scenario under consideration

In Step 3 of the base risk assessment, the ICT assets associated with an organisation have been modelled, and associated with the processes they support. In this first stage, the potential configurations of these assets are placed in a model of a potential cloud deployment. For the example video surveillance scenario, a potential organisation is presented in Figure 27. In the scenario, TenSys are considering migrating their anomaly detection and archiving services into the cloud, along with the associated data items, i.e., the archive and live video data. This arrangement could offer them a number of benefits, including a reduced on-site presence, the ability to more readily handle peak loads, e.g., when public safety incidents occur, and easier incorporation of new clients (i.e., adding new clients does not require significant capital expenditure in server equipment).

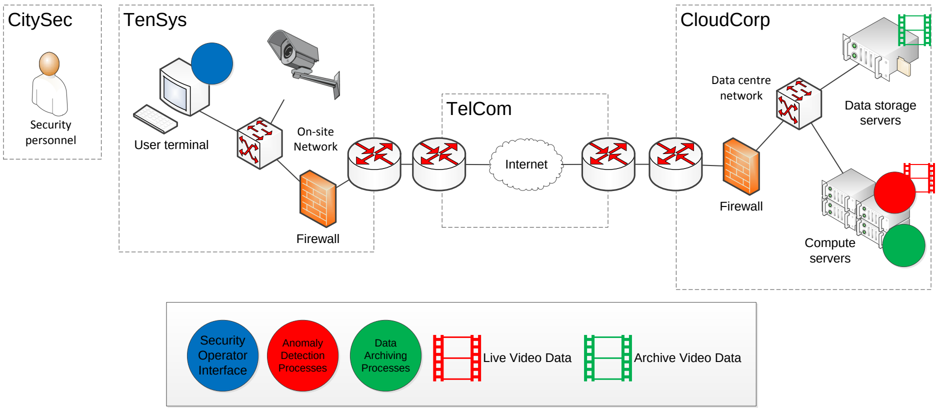

Figure 27: A potential deployment of the video surveillance ICT assets in the cloud

An important item to consider is the nature of the deployment model that is being offered, i.e., is this a private, public, community or hybrid cloud, and, if resource pooling (or shared tenancy) is part of the offering, what is the nature of the tenants the assets will be co-located with. This information can be used to examine the risks associated with the threats and vulnerabilities that are outlined in Section 6.1.3 on resource pooling. In general, an organisation should try to garner and model as much information as possible about the potential deployment that relates to the different threat and vulnerability categories, which are shown in Figure 12. In a later stage, these factors will be used to determine the risks for cloud adoption. One of the major changes that may occur through cloud deployment is a change in the trust boundaries of the components that interact – in the video surveillance, the separation of the `security operator interface` and the remaining system components in the cloud introduces a new trust boundary. (This may be the case within the cloud deployment context as well, if components are executed in different regions of the cloud data centre.) A widely-used approach to identifying such interactions is to develop a Data Flow Diagram (DFD), which forms part of the Microsoft STRIDE method [H+06]. An example DFD is shown in Figure 28.

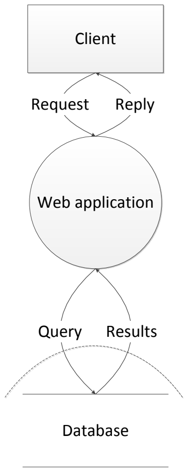

An example data flow diagram, showing a `Client` *input/output* to a `Web application` *function* or *process*, which interacts with the `Database` *data store*. The *data flows* are represented by the labelled arrows, e.g., `request/reply`, and the dashed line indicates there is a *trust boundary* between the `Web application` and the `Database`. We suggest a similar DFD should be produced for the system that is under consideration – the introduction of new trust boundaries may point to potential new threats and vulnerabilities, e.g., man-in-the-middle attacks.

Figure 28: An example Data Flow Diagram (DFD) that can be used to identify the trust boundaries between components in a system

Klöti *et al.* have combined the Microsoft STRIDE vulnerability analysis method, which includes the development of DFDs, and attack trees [Sch99] to carry out a security analysis of OpenFlow networks [MAB+08]. OpenFlow networks are being increasingly used to implement large-scale data centres, which support cloud computing services, with some commercial offering already available. Therefore, the issues raised by Klöti *et al.* regarding OpenFlow security will become increasingly relevant to understanding some of the security issues of cloud usage, if this trend persists.

Re-evaluate levels for existing risk scenarios and controls

Having modelled the potential cloud deployment, the next step is to re-examine the risk scenarios that were identified as part of the base risk assessment. In some cases, the existing risks will be reduced, and in other situations aggravated. We foresee the following possible changes that need to be made to the existing risk scenarios:

- *Risk scenarios are no longer valid*: it may be the case that entire risk scenarios are no longer relevant because of the introduction of cloud computing. If this is the case, they should be removed from the risk assessment model, or no longer associated with the assets that have been migrated to the cloud. However, we foresee this outcome being unlikely and the adoption of cloud computing will augment the risks in a more subtle way, as follows.

- *Modifications to the threats and vulnerabilities*: each risk scenario has a set of threats and vulnerabilities associated with it; these must be re-examined. Certain threats and vulnerabilities may no longer be relevant, and should be removed from the risk scenario. On a more detailed level, the severity and frequency of the vulnerabilities and threats, respectively, may have changed, and should be adjusted.

- *Changes to security controls that are in place*: as mentioned earlier, a number of security controls will be in place that affect the risks associated with the specified threats and vulnerabilities. Cloud deployment may make these controls more effective, or conversely less effective and more challenging to implement (or no longer required). The challenges associated with implementing security controls that are outlined in Section 6.1.10 may be helpful to support this analysis.

Model cloud risk scenarios, including threats and vulnerabilities

The use of cloud computing may result in entirely new risk scenarios, which should be considered as part of the risk assessment. These should be modelled as before for the base risk assessment; the process outlined in Figure 23 can be followed to include these new items. Additionally, using the cloud may introduce entirely new threats and vulnerabilities that must be considered, and added to the existing risk scenarios that were identified in the base risk assessment. To support this evaluation, we have created a comma separated value (CSV) file that contains the items in the catalogue described in Section 6. The CSV file can be downloaded from the SECCRIT website[14] and included in the Verinice tool using the process shown in Figure 29.

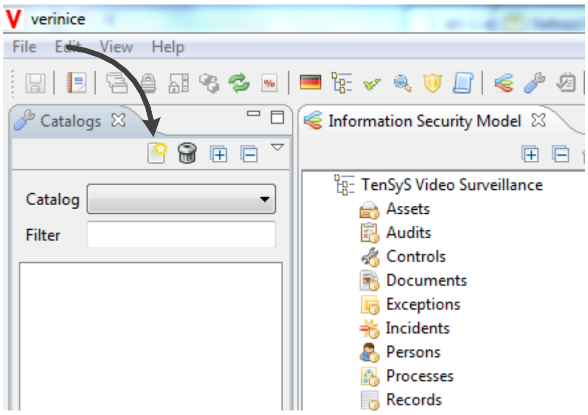

To import the SECCRIT vulnerability and threat catalogue that is in the CSV file format into Verinice, simply click on the button shown and open the file. The catalogue will appear in the area shown. Items can then be dragged and dropped into a suitable folder (threat or vulnerability) and subsequently associated with a risk scenario, as shown in Figure 23.

Figure 29: The SECCRIT threat and vulnerability catalogue can be imported into the Verinice tool.

Complete the cloud-based risk scenario

The final three stages of the SECCRIT risk assessment extension for cloud, i.e., steps 11 to 13 that are shown in Figure 19, are very similar in nature to steps that need to be carried out for the base risk assessment. As before the new risk scenarios must be associated with the assets that are augmented by the cloud usage (how to associate risk scenarios and assets is shown in Figure 24). Subsequently, the security protections that *could* be realised to mitigate

[14] `https://www.seccrit.eu`

the new risk scenarios, including threats and vulnerabilities, should be modelled and associated with the relevant scenarios and assets (following a similar process to that shown in Figure 25). Future work will investigate the protection measures that can be applied to the threats and vulnerabilities that are outlined in the SECCRIT catalogue. Subsequently, the new risk levels and reports should be calculated, using the process outlined in Figure 26.

Compare non-cloud and cloud scenario risks

Using the report generated by the Verinice tool, a comparison can be made between the different cloud and non-cloud deployment scenarios. For example, the overall risk situation for an organisation can be analysed, along with a more detailed examination of the specific risks to assets. Of course, the potential cyber-security risks are just one, all be it important, dimension to consider when an organisation considers migrating to the cloud, others include the cost savings and flexibility that cloud can afford.

7.3 Summary

In this section, we have presented an approach that organisations can use to understand the risks associated with migrating their assets to a cloud offering. In summary, the approach builds on the results from a standard asset-driven risk assessment, and augments it by modelling the potential cloud deployment, along with its effects on the risks, including threats and vulnerabilities, and security controls. We have indicated how the approach could be implemented via a video surveillance scenario, using the Verinice ISMS tool, which can be used to support ISO 27001 compliance. We argue that building on an existing risk assessment is reasonable, as organisations that support high-assurance ICT services (such as critical infrastructure providers) should be executing an information security risk assessment as part of their standard practices. This assumption will be strengthened in Europe as the proposed Network and Information Security (NIS) directive comes into force, which encourages risk assessments to be implemented by organisations in this sector, e.g., via the realisation of national legislation.

Future work in this area could include the development of suitable security protections that relate to the SECCRIT threat and vulnerability catalogue – for the moment, the organisation that is conducting the risk assessment has little support in this regard. Determining the severity of the vulnerabilities and the likelihood of threats occurring is one of the more challenging aspects of this risk assessment, especially with respect to the cloud. A suitable knowledge-base could be developed to support organisations with respect to determining these aspects. Furthermore, the concepts regarding measuring risk, which are discussed in Section 8, could support this activity.

8 Online Measurement of Risk in Cloud Environments

One of the significant difficulties when attempting to determine the risks associated with using cloud computing for high assurance ICT services is determining accurate, ideally quantitative, measures regarding the *probability* and *impact* of a particular threat. Acquiring reliable measures for these items is, in general, problematic for cyber-security risk assessment; however, the cloud makes this somewhat harder still. Characteristics of cloud computing, such as rapid elasticity, make it difficult to determine a stable representation of the infrastructure under scrutiny, e.g., with respect to shared tenancy and resource availability threats, and given that many cloud offerings (i.e., public and community clouds) involve multiple stakeholders, a lack of transparency can arise when third-parties need to evaluate the risks from others in the service-provisioning chain (see Section 6.1.7 on organisational issues).

These issues have lead some researchers to propose the concept of *Risk Assessment as a Service* (RAaaS) for the cloud – an on-demand service, wherein properties of the cloud infrastructure under inspection for risk assessment can be measured (by a cloud user) in a trustworthy manner [TTG13, KP10]. Such a service could be used as a basis for performing a continuous *online* risk assessment, which can mitigate the issues associated with the dynamic nature of cloud computing. Furthermore, enabling remote measurement of the cloud infrastructure by third-parties can enable the transparency required for an informed risk assessment. The concept of RAaaS is arguably still in its infancy, with a number of open research issues that need to be addressed, including determining the trustworthiness of the data that is measured, addressing potential privacy (or commercial sensitivity) issues in shared tenancy settings, and instrumenting an infrastructure in a resource-efficient manner. A notable open issue is determining *specifically* what it is that should be measured in the cloud infrastructure that can support a risk assessment – it is this issue we aim to take initial steps to address here, using the SECCRIT threat and vulnerability catalogue from Section 6.

In order to make suggestions about what should be measured as part of a potential RAaaS, it is first necessary to revisit how risk is determined. There are several definitions of risk, but typically the formulation presented in Equation 1 is used, i.e., *risk* is the product of the potential *impact* of a *threat* and an associated *vulnerability*.

$$risk = impact \times threat(probability) \times vulnerability(severity) \tag{1}$$

The impact associated with a risk can be measured in a number of ways, including reduced organisational reputation (leading to reduced future business) or direct monetary loss associated with an incident, e.g., due to penalties from regulators or those imposed by Service Level Agreement (SLA) failures. In any case, the potential impact of a risk is very much service- or industry-specific, and is beyond the scope of this discussion. Meanwhile, the *probability* of a threat occurring can be determined in a number of ways, including using advisories, such as the CSA's Top Threat Catalogue [CSA13] or ENISA's telecommunications incident report [Dek12]. Furthermore, ongoing research is investigating how online measures, e.g., from intrusion detection systems, can be used to augment threat probability indicators [ASH$^+$05]. Finally, measuring properties of a vulnerability on risk can incorporate a number of factors, including the availability of usable exploits and structural properties of an infrastructure, e.g, the presence of redundant components, longevity of standby electricity generators or the extent networks are isolated in a cloud network infrastructure. In many cases, the determined values for impact, threat probability and vulnerability severity are normalised into the range $[0, 1]$. These three factors are intimately related: the probability of a cybersecurity threat manifesting can be affected by the nature of a vulnerability, and in some cases, the potential impact it could have, e.g., relating to terrorist of hacktivist actors.

Therefore, relating this discussion back to RAaaS, our aim is to specifically identify the *measurable properties* of a cloud infrastructure that can support the risk analyst estimate the *probability* of a threat manifesting or determine the *severity* of a vulnerability, such as its ability to be exploited. We achieve this by examining each of the threats and vulnerabilities that are outlined in the SECCRIT catalogue, and when possible identify cloud infrastructure properties that could be used to determine the aforementioned factors. To support this analysis, we used the SECCRIT architectural model, which is summarised in Figure 10 on page 28.

In SECCRIT deliverable D5.1 [Rol13], we have described the SECCRIT architectural framework and various views on it. The monitoring view shown in Figure 30 illustrates which properties of the system can be measured on which layer. At the *Cloud Infrastructure* layer, the resource utilization of all physical resource, e.g., CPU load, network congestion, or storage utilization, is available as well as hardware specific measurements, e.g., CPU temperature, electrical grid utilization, or system maintenance activities. A subset of these measures is exposed to the tenants, and is usually limited to a tenant-specific view – an exterior view on the virtual resources at the *Tenant Infrastructure* layer. In addition, an interior view on the virtual resource by using measurement agents within the virtual environment complements the measured properties at this layer. Service-specific measurements have to be done at the *Service* layer, e.g., average response time, service availability or maximum throughput. Finally, the service users observe the delivered quality of service (QoS) and check this against the agreed SLA.

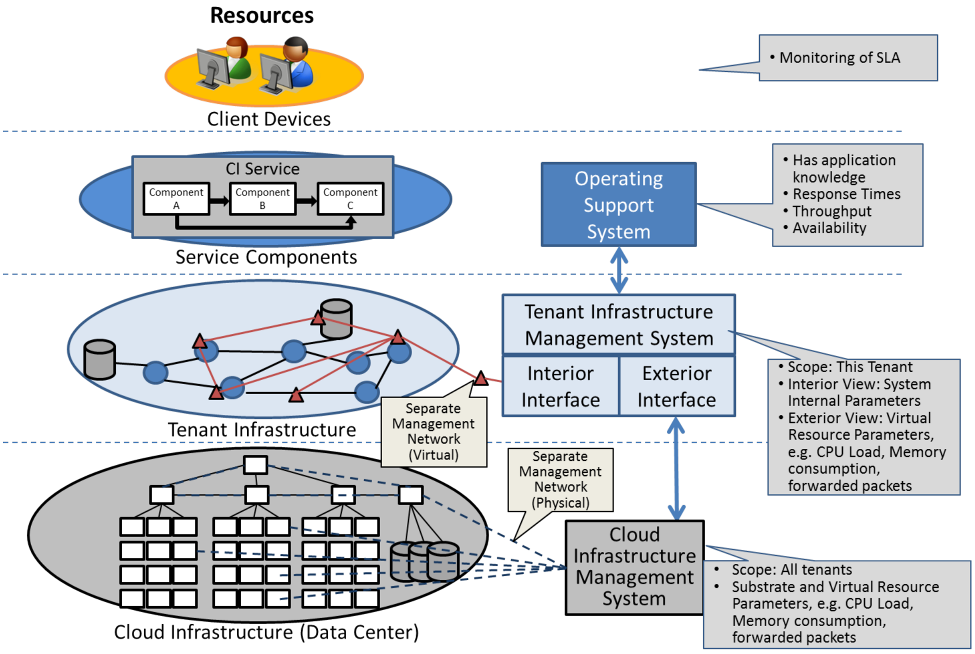
Resources
• Monitoring of SLA
Client Devices
CI Service
Component A
Component B
Component C
Service Components
Operating Support System
• Has application knowledge
• Response Times
• Throughput
• Availability
Tenant Infrastructure Management System
Interior Interface
Exterior Interface
• Scope: This Tenant
• Interior View: System Internal Parameters
• Exterior View: Virtual Resource Parameters, e.g. CPU Load, Memory consumption, forwarded packets
Separate Management Network (Virtual)
Tenant Infrastructure
Separate Management Network (Physical)
Cloud Infrastructure Management System
• Scope: All tenants
• Substrate and Virtual Resource Parameters, e.g. CPU Load, Memory consumption, forwarded packets
Cloud Infrastructure (Data Center)

Figure 30: The SECCRIT monitoring architecture

The upper two layers of Figure 30 are unmodified with respect to legacy service deployments, and the respective metrics are well-understood and documented. Thus, we will focus on the infrastructure related metrics – both of physical and virtual resources. A cloud provider uses a plurality of subsystems to monitor the physical resources. Many of these measurements will not be exposed to the tenants, e.g., over-subscription ratio of physical resources. Such measures will be *transparent* to the tenants. But tenants running critical services will ask for an increased *translucency* of such measurements, in order to perform an online risk assessment. The degree of *translucency* will be subject to business models and market demands; critical infrastructure providers might prefer contracting with cloud providers which provide the required translucency. Thus, we focus on the technical aspects and provide a list of measurable properties, which can help an online risk assessment.

OpenStack, as one of the most relevant open-source cloud platforms, provides two modules for monitoring, which are called Ceilometer[15] and Healthnmon[16]. These general monitoring frameworks provide a large set of metrics at the *Cloud Infrastructure* level – similar modules

[15] https://wiki.openstack.org/wiki/Ceilometer
[16] https://wiki.openstack.org/wiki/CloudInventoryManager

are available for competing cloud management solutions. The set of metrics defined in Table 14 is considered helpful for a risk assessment, which can be provided by the cloud infrastructure provider to the tenant infrastructure provider. On the *Tenant Infrastructure Provider* layer, a large set of measures can be used in a risk assessment process. A number of these are summaried in Table 15.

Table 14: The metrics that can be provided by the cloud infrastructure to its tenants, in order to support online risk assessment

Metric	Description	Cat. ID
Node-ID [node-id]	The ID of the compute node to which a virtual machine is mapped can be used to verify that redundant service components are mapped to independent physical resources. The ID might be a tenant-specific identifier to preserve data centre secrets.	OI-1
Physical component failure [node-id]	A service component which requested to be scheduled to a compute node [node-id] with certain hardware features, e.g., two or more redundant network interface cards, should be informed about failures of such hardware to initiate a fail-over. Until the VM is migrated to a new compute node with the requested hardware features the service is at higher risk.	VN-6, PI-2
Time to resource migration [s]	Tenants scheduled to go into maintenance mode will be informed about planned or impending migrations of virtual resources ahead of time. The tenant can take preventive measures to reduce the risk of impact on the service delivered, e.g., force a fail-over during low service utilisation before migration is performed.	RE-10, RE-11
Attack notification [node-id]	If the cloud provider detects an active attack and takes appropriate counter-measures, the tenant needs to be informed about these measures. The cloud provider will provide a notification for each compute node affected. Examples can be firewall rules imposed.	RP-7, MS-3

The measurable properties of a cloud infrastructure, which are enumerated in Tables 14 and 15, represent an early stage in developing the RAaaS concept. Future work in the SECCRIT project will contribute on various aspects to this concept. In work package four, the work on anomaly detection on the network as well as on the virtual machine level will provide measurements on the cloud infrastructure layer. Furthermore, the work on a cloud resilience framework will build closed control loops based on these measurements. The activities conducted in work package five on a trustworthy audit trail and validation infrastructure that can support root cause analysis will be used to move this concept further.

9 Risk Management and Cloud Computing

It is a basic observation that risk is relational [Gar03, BC11]. Things become risks when they have the potential to harm our interests, and it is our interests that shape what kind of risks they are. Risk is also normative: when we create risks we create social obligations to mitigate them

Table 15: The metrics that are available at the tenant infrastructure level

Metric	Description	Cat. ID
Virtual CPU load [%]	This metric relates to the utilization of the virtual compute resources, which could be measured to indicate that a virtual service is behaving anomalous, e.g., with respect to its normal behaviour patterns, or is under-resourced.	RE-1, RE-4
Virtual memory load [%]	In a similar fashion to virtual CPU load, unusual or undesirable virtual memory usage could be indicative of a problem, e.g., a misconfiguration or attack behaviour, which could leas to a service failing, for example.	RE-1, RE-4
Virtual link load [%]	Unusual utilization of the virtual link capacity could be indicative of threats, such as a Distributed Denial of Service attack (originating or targeted at a virtual service). Furthermore, if malicious behaviour is not present, over utilisation of virtual links could point to a misconfiguration or resource over-utilisation issue.	RE-2, RE-5
Number of packets [#]	Related to the virtual link load, the number of packets that are observable at the tenant infrastructure level could prove useful for risk assessment. This metric offers a subtlety different view than the link load, as an unusual number of small packets could be sent, which do not result in high link utilisation.	RE-2, RE-5
Number of active connections [#]	Another related network metric is the average number of packets that have been forwarded within an interval. Again, this metric could be used to determine the existence of network-borne threats or those that result in unusual network behaviour	RE-2, RE-5
Number of blocked packets [#]	This is a firewall statistic that can be used to determine the number of packets that have triggered a firewall rule, which resulted in it being dropped. Such dropped packets could indicate (the level of) malicious behaviour, and point to increased risks from cyber-attacks.	
Virtual switch congestion [s]	The virtual machine instance are connected via a virtual switch; this metrics relates to the average packet delay that is introduced by this switch. As with the other network characteristics, anomalous values of this measure could indicate problems.	RE-2, Re-5
Virtual storage load [%]	This metric indicates the utilisation of the virtual storage – anomalous usage of the virtual storage, e.g., a rapid growth in the storage utilisation, could point to malicious behaviour or misconfigurations, for example.	RE-3, RE-6

or contain them [Luh93, Dou92, Sch99]. Thus it is actors' failure to live up to their obligations rather than failure to manage risk down to certain levels per se, that often causes most concern [Fre03]. And risk is a social construction. Hilgartner [Hil92] suggests that any technology is a whole network of things, including physical entities but also social processes and institutions that make them work. Within such a network, particular entities become prominent as risk objects, especially when we can trace a clear link between them and some harm, or when they appear resistant to control. In the process of determining what the risk objects are we distribute responsibilities in a particular way that serves or harms the interests of different parties.

This suggests some important points about risk and cloud computing. First, it is necessary to understand the interests of the relevant actors, and user organisations in particular. In education and public services these interests are often related to integrity and reputation. User organisations are likely to deal in personal data about other parties, such as student grades and reports, or client problems and conditions. Such organisations have to be seen to protect such data, and meet both moral and legal obligations in relation to maintaining confidentiality and integrity. Putting data into the cloud presents an obvious and basic threat and naturally creates the perception of a vulnerability. Risk management needs to start from an analysis of such issues.

Second, risks have to be traced not only to their causes but to those given responsibility for controlling them. Risk management strategies have to be capable of letting such actors show that they are discharging their obligations. Legal agreements may be important, as may be warnings to and restrictions on end users. Given the way cloud computing distributes obligations across user and provider organisations, and possibly to others acting as sub-contractors to the provider, it becomes more important for the limits of responsibility to be clear. It may be difficult for a user organisation to constrain cloud use by its organisational members. Furthermore, it may also be difficult for the organisation to monitor the actions of a provider with whom it cannot engage at a managerial level. And it may be difficult for the organisation to demonstrate to its clients or the public that it takes such risks seriously without simply avoiding the use of cloud services entirely.

Third, there should be a clear understanding of the network of objects that make up the technology. This includes all the provider's technical systems and organisational processes, the technical systems that remain within user organisations, communication systems, security systems and processes, users, user managers, provider system managers, malign third parties of various kinds, and so on. The network is larger for cloud computing than conventional computing, and less visible to user organisations in particular. Different actors will probably stress different entities as risk objects, and it is important to look at the implications of any discrepancies. The migration of computing from one physical site to another, for example, might be seen as a risk mitigating strategy by a provider but a potential source of risk by users – especially if the geographical location of data is relevant to their interests (for example, if the relevant laws vary by region). Different actors may stress certain risk objects in a way that distributes responsibility away from themselves, which could leave significant gaps in defences. Users might see risks of sharing tenancy with a highly targeted other user as naturally being the responsibility of the provider, but providers might evade such responsibility in their contractual provisions.

10 Conclusions and Further Work

Understanding the risks associated with deploying high-assurance ICT services in the cloud is of critical importance, as these services support the critical infrastructures that our society depends on. Without understanding these risks, it is not clear whether the cloud can be used securely and safely, and in which forms it could be applied to support these services. To this end, this deliverable has provided a number of items that can be used by the organisations that implement high-assurance ICT services to understand cloud-based risks.

Central to this is a novel cloud-specific threat and vulnerability catalogue, which can be be used to support a risk assessment. We have shown how this catalogue can be applied to determine the risk associated with adopting the cloud, through the use of an extension to existing risk assessment processes. We have performed a study of risk perceptions, with respect to the cloud, from an individual and organisational perspective. This study suggests the threats perceived by individuals can be wide-ranging, and a major concern of a university organisation – the target of our organisational analysis – is inappropriate use of its data by the cloud provider, such as its dissemination to third-parties. Determining the likelihood of threats emerging and the severity of vulnerabilities is challenging for the cloud, because of a lack of transparency and their highly dynamic nature. As a starting point to address this issue, we have pointed to some initial metrics that could be measured in a cloud infrastructure to support this analysis in an online manner.

There is a great deal of future work that could be done in this area. Whilst the threat and vulnerability catalogue we propose is arguably extensive, further work is needed to identify the *controls* that can be applied to mitigate these. This is both important practically, i.e., to be able to address the threats, and to support the implementation of a risk assessment – these controls augment the likelihood and impact indicators in an assessment. Our efforts toward being able to conduct online measurements of risk for a cloud deployment are preliminary, with many open questions needing to be addressed. For example, it is not clear how some of these measures can be determined in a trustworthy manner, i.e., ensured they have not been subjected to tampering. Furthermore, in some cases, data could be revealed to tenants in an infrastructure that could be commercially sensitive or used as the basis for an attack – an analysis of these threats is required.

References

[Ada13] Kirila Adamova. Anomaly Detection with Virtual Service Migration in Cloud Infrastructures. Master's thesis, ETH Zürich, March 2013.

[Alu12] Alura Support. How Hurricane Sandy Forced IT Professionals to Rethink Cloud Computing. http://alurasolutions.com/how-hurricane-sandy-forced-it-professionals-to-rethink-cloud-computi November 2012.

[Ama13] Amazon. Amazon Elastic Compute Cloud: Regions and Availability Zones. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html, October 2013. User Guide (API Version 2013-10-01).

[ANM11] I.M. Abbadi, C. Namiluko, and A. Martin. Insiders analysis in Cloud computing focusing on home healthcare system. In *International Conference for Internet Technology and Secured Transactions (ICITST)*, pages 350–357, Abu Dhabi, United Arab Emirates, December 2011.

[ASH^{+}05] André Arnes, Karin Sallhammar, Kjetil Haslum, Tonnes Brekne, Marie Moe, and SveinJohan Knapskog. Real-Time Risk Assessment with Network Sensors and Intrusion Detection Systems. In *Computational Intelligence and Security (CIS 2005)*, volume 3802, pages 388–397, Xian, China, December 2005.

[Aug09] R. Auger. SQL Injection. http://projects.webappsec.org/SQL-Injection, 2009.

[Bar12] David Barker. A Guide to Physical Security for Data Centers. http://www.datacenterjournal.com/facilities/a-guide-to-physical-security-for-data-centers/, July 2012.

[BC11] A. Boholm and H. Corvellec. A relational theory of risk. *Journal of Risk Research*, (14):175–190, 2011.

[BDP06] Stefano Bistarelli, Marco Dall'Aglio, and Pamela Peretti. Strategic Games on Defense Trees. In *Formal Aspects in Security and Trust (FAST)*, pages 1–15, Hamilton, Ontario, Canada, August 2006.

[BFMR10] K. Butler, T.R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1):100–122, 2010.

[BHC13] Roland Burger, Elmar Husmann, and Christian Cachin. Cloud, Trust, Privacy: Trustworthy cloud computing whitepaper. http://www.tclouds-project.eu/index.php/published-results, 2013.

[Blo11] Jason Bloomberg. Cloud brokering: Building a cloud of clouds. http://www.zdnet.com/blog/gardner/cloud-brokering-building-a-cloud-of-clouds/4140, April 2011.

[BMP^{+}13] S. Bleikertz, T. Mastelic, S. Pape, W. Pieters, and T. Dimkov. Defining the Cloud Battlefield – Supporting Security Assessments by Cloud Customers. In *IEEE International Conference on Cloud Engineering (IC2E)*, pages 78–87, San Francisco, CA, USA, March 2013.

[BP13] Silvia Balaban and Frank Pallas. SECCRIT Deliverable D2.2 Legal fundamentals. https://seccrit.eu/publications/publicreports, June 2013.

[BRC12] Armaghan Behnia, Rafhana Rashid, and Junaid Chaudhry. A Survey of Information Security Risk Analysis Methods. *Smart Computing Review*, 2(1), February 2012.

[BWT12] Salman A. Baset, Long Wang, and Chunqiang Tang. Towards an understanding of oversubscription in cloud. In *Proceedings of the 2nd USENIX conference on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services*, Hot-ICE'12, pages 7–7, Berkeley, CA, USA, 2012. USENIX Association.

[CCR09] F. Callegati, W. Cerroni, and M. Ramilli. Man-in-the-Middle Attack to the HTTPS Protocol. *IEEE Security & Privacy*, 7(1):78–81, 2009.

[CFH+05] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul, Christian Limpach, Ian Pratt, and Andrew Warfield. Live migration of virtual machines. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 273–286, 2005.

[CGC06] Francisco L. Crespo, Miguel A.A. Gómez, and Javier Candau. *MAGERIT – version 2 Methodology for Information Systems Risk Analysis and Management Book I - The Method*. Ministerio de administraciones públicas, 2006.

[cS13] Egemen K. Çetinkaya and James P. G. Sterbenz. A Taxonomy of Network Challenges. In *Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN)*, pages 322–330, Budapest, March 2013.

[CSA13] CSA Top Threats Working Group. The Notorious Nine: Cloud Computing Top Threats in 2013. Available on-line at: https://cloudsecurityalliance.org/research/top-threats/, February 2013. Cloud Security Alliance (CSA).

[CXZB11] Ashley Chonka, Yang Xiang, Wanlei Zhou, and Alessio Bonti. Cloud security defence to protect cloud computing against HTTP-DoS and XML-DoS attacks. *Journal of Network and Computer Applications*, 34(4):1097–1107, 2011.

[dBDG+03] Folker den Braber, Theo Dimitrakos, Bjørn Axel Gran, Mass Soldal Lund, Ketil Stølen, and Jan Øyvind Aagedal. The coras methodology: model-based risk assessment using uml and up. *UML and the Unified Process*, pages 332–357, 2003.

[Dek12] M. Dekker. Critical Cloud Computing: A CIIP perspective on cloud computing services. ENISA report, available online: http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/critical-cloud-computing, December 2012.

[DKL13] Marnix Dekker, Christoffer Karsberg, and Matina Lakka. Annual Incident Reports 2012: Analysis of Article 13a annual incident reports. http://www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/annual-reports/annual-incident-reports-2012, August 2013. ENISA.

[Dou92] M. Douglas. *Risk and Blame*. London: Routledge, 1992.

[ENI12] ENISA. Cloud Computing: Benefits, risks and recommendations for information security. https://resilience.enisa.europa.eu/cloud-security-and-resilience/publications/cloud-computing-benefits-risks-and-recommendations-for-information-se December 2012. Editors: Lionel Dupré and Thomas Haeberlen.

[ENI13] ENISA. Inventory of risk management/risk assessment methods and tools. http://www.enisa.europa.eu/activities/risk-management/current-risk/risk-management-inventory, 2013.

[Eur13] European Commission. Proposal for a DIRECTIVE OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL concerning measures to ensure a high common level of network and information security across the Union. http://ec.europa.eu/digital-agenda/en/news/eu-cybersecurity-plan-protect-open-internet-and-online-freedom-and-o 2013. COM(2013) 48 final.

[FFCdM11] Andreas Fischer, Ali Fessi, Georg Carle, and Hermann de Meer. Wide-Area Virtual Machine Migration as Resilience Mechanism. In *Proceedings of the 2011 IEEE 30th Symposium on Reliable Distributed Systems Workshops*, SRDSW '11, pages 72–77, Washington, DC, USA, 2011. IEEE Computer Society.

[Fre03] W.R. Freudenberg. Institutional failure and the organizational amplification of risk: The need for a closer look. In *The Social Amplification of Risk*, pages 102–120. Cambridge University Press, 2003.

[FvCvEW10] V.N.L. Franqueira, A. van Cleeff, P. van Eck, and R. Wieringa. External insider threat: A real security challenge in enterprise value webs. In *ARES '10 International Conference on Availability, Reliability, and Security, 2010*, pages 446–453, Krakow, Poland, February 2010.

[Gar03] D. Garland. The rise of risk. In R.V. Ericson and A. Doyle, editors, *Risk and Morality*, pages 48–86. University of Toronto Press (Toronto), 2003.

[GWS11] B. Grobauer, T. Walloschek, and E. Stocker. Understanding cloud computing vulnerabilities. *IEEE Security Privacy*, 9(2):50–57, 2011.

[H$^+$06] Shawn Hernan et al. Uncover Security Design Flaws Using The STRIDE Approach. http://msdn.microsoft.com/en-gb/magazine/cc163519.aspx, 2006.

[HGX08] Zhongsheng Hua, Bengang Gong, and Xiaoyan Xu. A DS-AHP approach for multi-attribute decision making problem with incomplete information . *Expert Systems with Applications*, 34(3):2221–2227, 2008.

[Hig09] Kelly Higgins. Hacking Tool Lets A VM Break Out And Attack Its Host. http://www.darkreading.com/applications/hacking-tool-lets-a-vm-break-out-and-att/217701908, June 2009.

[Hil92] S. Hilgartner. The social construction of risk objects: Or, how to pry open networks of risk. In J.F. Short and L. Clarke, editors, *Organizations, Uncertainties, and Risk*, pages 39–53. Westview Press (Boulder, CO), 1992.

[Int13] International Society of Automation (ISA). ISA99, Industrial Automation and Control Systems Security. http://www.isa.org/MSTemplate.cfm?MicrositeID=988&CommitteeID=6821, 2013.

[JKKD12] Ming Jiang, Mariam Kiran, Tom Kirkham, and Karim Djemame. D6.1.3.2 Scientific Report on Risk Assessment in OPTIMIS. Available online at: http://www.optimis-project.eu/sites/default/files/content-files/document/optimis-public-delievarable-d6132-scientific-report-risk-assessment.pdf, May 2012. OPTIMIS project deliverable.

[Joi13] Joint Task Force Transformation Initiative. Security and Privacy Controls for Federal Information Systems and Organizations. http://dx.doi.org/10.6028/NIST.SP.800-53r4, April 2013. NIST Special Publication 800-53 Revision 4.

[KKS13] R. Klöti, V. Kotronis, and P. Smith. OpenFlow: A Security Analysis. In *8th Workshop on Secure Network Protocols (NPSec 2013)*, Göttingen, Germany, October 2013.

[KP10] Burton S. Kaliski, Jr. and Wayne Pauley. Toward risk assessment as a service in cloud environments. In *2nd USENIX conference on Hot topics in cloud computing*, pages 13–13, Berkeley, CA, USA, June 2010.

[KS05] Bilge Karabacak and Ibrahim Sogukpinar. Isram: information security risk analysis method. *Computers & Security*, 24(2):147–159, 2005.

[Lem12] Robert Lemos. Insecure API Implementations Threaten Cloud. http://www.darkreading.com/cloud/insecure-api-implementations-threaten-cl/232900809, April 2012.

[LET+11] Rafal Leszczyna, Elyoenai Egozcue, Luis Tarrafeta, Victor F. Villar, Ricardo Estremera, and Jairo Alonso. Protecting Industrial Control Systems Recommendations for Europe and Member States. http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/scada-industrial-control-systems/protecting-industrial-control-syste-recommendations-for-europe-and-member-states/at_download/fullReport, 2011. ENISA whitepaper.

[LT75] Harold A Linstone and Murray Turoff. *The delphi method*. Addison-Wesley Reading, MA, 1975.

[Luh93] N. Luhmann. *Risk: A Sociological Theory*. Berlin: Walter de Gruyter, 1993.

[MAB+08] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: Enabling Innovation in Campus Networks. *SIGCOMM Comput. Commun. Rev.*, 38(2):69–74, March 2008.

[MG11] Peter Mell and Timothy Grance. The NIST Definition of Cloud Computing. Technical Report Special Publication 800-145, National Institute of Standards and Technology (NIST), September 2011.

[MR04] Jelena Mirkovic and Peter Reiher. A taxonomy of DDoS attack and DDoS defense mechanisms. *SIGCOMM Computer Communications Review*, 34(2):39–53, April 2004.

[MS10] David Molnar and Stuart E. Schechter. Self Hosting vs. Cloud Hosting: Accounting for the Security Impact of Hosting in the Cloud. In *Workshop on the Economics of Information Security (WEIS)*, Cambridge, MA, USA, June 2010.

[MSWA03] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. User-level internet path diagnosis. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*, SOSP '03, pages 106–119, New York, NY, USA, 2003. ACM.

[Nei12] Mike Neil. Windows Azure Service Interruption in Western Europe Resolved, Root Cause Analysis Coming Soon. `http://blogs.msdn.com/b/windowsazure/archive/2012/07/27/windows-azure-service-interruption-in-western-europe-resolved-root-c` `aspx`, July 2012. MSDN Blog Post.

[Nor13] North American Electric Reliability Corporation (NERC). CIP Standards. `http://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx`, 2013.

[PJGLSAH11] Lluis Pamies-Juarez, Pedro García-López, Marc Sánchez-Artigas, and Blas Herrera. Towards the design of optimal data redundancy schemes for heterogeneous cloud storage infrastructures. *Computer Networks*, 55(5):1100–1113, 2011.

[PMS13] N Pramod, Anil Kumar Muppalla, and KG Srinivasa. Limitations and Challenges in Cloud-Based Applications Development. In *Software Engineering Frameworks for the Cloud Computing Paradigm*, pages 55–75. Springer, 2013.

[Ree10] Jeff Reed. Following Incidents into the Cloud. `http://www.sans.org/reading-room/whitepapers/incident/incidents-cloud-33619`, September 2010. SANS Institute.

[Rol13] Roland Bless, et.al. Deliverable 5.1 design and api for audit trails and root-cause analysis, 2013.

[Row13] Rowan Klöti. OpenFlow: A Security Analysis. MSc thesis, D-ITET, ETH Zurich. `ftp://ftp.tik.ee.ethz.ch/pub/students/2012-HS/MA-2012-20.pdf`, 2013.

[Sch99] B. Schneier. Attack Trees. *Dr. Dobb's Journal*, December 1999.

[SFS11] Keith Stouffer, Joe Falco, and Karen Scarfone. Guide to Industrial Control Systems (ICS) Security. `http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf`, June 2011. NIST Special Publication 800-82.

[SHc$^+$10] James P. G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks: Special Issue on Resilient and Survivable Networks (COMNET)*, 54(8):1245–1265, June 2010.

[SK11] S. Subashini and V. Kavitha. A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, 34(1):1–11, 2011.

[SKGK10] Alan Shieh, Srikanth Kandula, Albert Greenberg, and Changhoon Kim. Seawall: performance isolation for cloud datacenter networks. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, HotCloud'10, pages 1–1, Berkeley, CA, USA, 2010. USENIX Association.

[Sof13] Software Engineering Institute Carnegie Mellon. OCTAVE Information Security Risk Evaluation. http://www.cert.org/octave/, 2013.

[ST12] Lwin Khin Shar and Hee Beng Kuan Tan. Defending against Cross-Site Scripting Attacks. *Computer*, 45(3):55–62, 2012.

[SuhSHB13] Steven Simpson, Noor ul-hassan Shirazi, David Hutchison, and Helge Backhaus. Anomaly detection techniques for cloud computing. https://www.seccrit.eu, 2013. SECCRIT Deliverable D3.1.

[SW10] P. Saripalli and B. Walters. QUIRC: A Quantitative Impact and Risk Assessment Framework for Cloud Security. In *IEEE 3rd International Conference on Cloud Computing (CLOUD)*, pages 280–288, Miami, FL, USA, July 2010.

[SWYP12] F.T. Sheldon, John Mark Weber, Seong-Moo Yoo, and W. David Pan. The Insecurity of Wireless Networks. *IEEE Security & Privacy*, 10(4):54–61, 2012.

[Tal13] Chris Talbot. Cloud Outages: Power Loss Blamed as Main Cause. http://talkincloud.com/cloud-computing-research/cloud-outages-power-loss-blamed-main-cause, March 2013.

[THGL11] Mark Taylor, John Haggerty, David Gresty, and David Lamb. Forensic investigation of cloud computing systems. *Network Security*, 2011(3):4–10, 2011.

[Tie99] K.J. Tierney. Toward a critical sociology of risk. *Sociological Forum*, (14):215–242, 1999.

[TJA10] H. Takabi, J.B.D. Joshi, and Gail-Joon Ahn. Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security & Privacy*, 8(6):24–31, 2010.

[TTG13] Marianthi Theoharidou, Nikolaos Tsalis, and Dimitris Gritzalis. In Cloud We Trust: Risk-Assessment-as-a-Service. In *7th IFIP WG 11.11 International Conference on Trust Management*, pages 100–110, Malaga, Spain, June 2013.

[Vek13] Nicos Vekiarides. How to Protect Against Your Cloud Storage Provider's Demise. http://www.drj.com/articles/online-exclusive/how-to-protect-against-your-cloud-storage-providers-demise.html, October 2013.

[Ver13] Jason Verge. Network Issues Cause Amazon Cloud Outage. http://www.datacenterknowledge.com/archives/2013/09/13/network-issues-cause-amazon-cloud-outage/, September 2013.

[WL03] D. Welch and S. Lathrop. Wireless security threat taxonomy. In *IEEE Systems, Man and Cybernetics Society Information Assurance Workshop, 2003*, pages 76–83, 2003.

[WLK^{+}12] Ping Wang, Wen-Hui Lin, Pu-Tsun Kuo, Hui-Tang Lin, and Tzu Chia Wang. Threat risk analysis for cloud security based on Attack-Defense Trees. In *8th International Conference on Computing Technology and Information Management (ICCM)*, pages 106–111, Seoul, Korea (South), April 2012.

[ZJRR12] Yinqian Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. Cross-vm side channels and their use to extract private keys. In *Proceedings of the 2012 ACM conference on Computer and communications security*, CCS '12, pages 305–316, Raleigh, North Carolina, USA, 2012.

A Legal Considerations

This appendix summarizes the legal issues that have been addressed in the preparation of this deliverable.

1. *Does the mechanism (or the system/context in which the mechanism is to be used) collect, process and/or store data that is attributed to or can be reattributed to natural persons ("personal data" as defined in the deliverable on legal fundamentals)?*

   No. There should be no need for the use of personal data for the purposes of the risk assessment approaches that are outlined in this deliverable.

2. *If yes:*

   (a) *for what reasons/purpose is the personal data collected etc.?*

   Not applicable.

   (b) *What mechanisms have been applied to minimize the amount of personal data being collected, processed, stored etc.? In particular, this can also include pseudonymization and anonymization techniques.*

   Not applicable.

3. *How does the mechanism ensure that non-authorized third parties do not have access to the data? And in the cases they did have, how can this be recognized and revealed? How are availability and integrity of the personal data ensured technically?*

   Not applicable.

4. *How can it be retraced in real time where the data are located (including country, data center, virtual machine, physical machine) and which processes/entities/etc. accessed and processed the data (so that cloud behavior becomes as transparent as possible to the cloud user and possibly other stakeholders)?*

   Not applicable.

5. *How can it be ensured and technically proven that, when legal givens require data to be deleted, the data are deleted everywhere?*

   Not applicable.

6. *Can nonfunctional requirements be configured into the technical system like: no outsourcing to a foreign country, delete data after 3 weeks, no co-location of datasets X and Y, . . . ?*

   Not applicable.

7. *In cases of a failure happening during runtime, how can it be retraced where the fault actually happened and what had been the cause of the failure?*

   Not applicable.

8. *What technical mechanisms are employed to make this "digital evidence" credible (trusted, not manipulated, provably true content, . . .) and to make it available to legitimate parties? (think of, e.g., a cloud user having to proof that a cloud provider did not fulfill his obligations)*

   Not applicable.

B The SECCRIT Risk Perceptions Questionnaire

The questionnaire that was used to conduct the individual cloud computing risk perceptions survey, which is described in Section 5.

Perceptions of Risk in Cloud Computing for Critical Infrastructures: A

Introduction

SECCRIT - SEcure Cloud computing for CRitical Infrastructure IT, is conducting this survey for perceptions of risk in cloud computing for critical infrastructure.

The purpose of this questionnaire is to gather the experiences and views of informed people on the risks of cloud computing, especially when used by organizations that provide critical infrastructure services. It is being used as part of the SECCRIT project, in order to find out what risks are seen as being particularly important by both users and providers of cloud computing.

Why should I complete this questionnaire?

1. It is important to learn about the experiences and viewpoints of users and providers to produce effective ways of managing risks.
2. We hope that you would find the questionnaire thought provoking.
3. If you choose to provide contact details you may find it informative to see what other people think about the risks of cloud services.

The results will be analysed and they may be used in scientific publications. These publications will present only a general analysis of your responses, and will not identify any people or organizations responding to the survey. All the responses you give will be held by us in strict confidence, and will be stored without any information that could identify you or your organizations.

Our Contact Details:
Noorulhassan Shirazi, n.shirazi@lancaster.ac.uk
Jerry Busby, busbyj@exchange.lancs.ac.uk

1. We would like to keep a separate record of your contact details so we can contact you again in a follow-up survey, but if you would prefer not to give your details you do not need to. If you give your contact details we will send you a summary of the findings. If you are happy to provide your email, please enter it below.

Perceptions of Risk in Cloud Computing for Critical Infrastructures: A

Information about you

2. Is your organization a provider, user or possible future user of cloud services?

The term cloud services is qualified as any of the three main cloud service models: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) or Software-as-a-Service (SaaS).

○ Provider

○ User

○ Possible future user

Other (please specify)

3. What is the principal industry of your organization?

***4. How many employees are there in the organization that you are affiliated to?**

○ 1-9 employees.

○ 10-50 employees.

○ 51-250 employees.

○ Over 250 employees.

***5. What is your role in this organization?**

6. How long have you been in this or similar roles?

○ Less than one year

○ 1-2 years

○ 3-6 years

○ More than 6 years

7. How many years of PERSONAL experience do you have in using, providing, or otherwise being engaged in cloud computing of the following kinds?

	No experience	Less than a year	1-2 Years	3-6 Years	More than 6 Years
Infrastructure-as-service (where the cloud provider provides a virtual machine and operating system)	○	○	○	○	○
Platform-as-service (providing an application programming interface)	○	○	○	○	○
Software-as-service (for example email, file sharing and so on)	○	○	○	○	○

8. How many years of experience does your current ORGANIZATION have in using, providing, or otherwise being engaged in cloud computing of the following kinds?

	No experience	Less than a year	1-2 Years	3-6 Years	More than 6 Years
Infrastructure-as-service (where the cloud provider provides a virtual machine and operating system)	○	○	○	○	○
Platform-as-service (providing an application programming interface)	○	○	○	○	○
Software-as-service (for example email, file sharing and so on)	○	○	○	○	○

9. If you are not currently a user of cloud services, but a potential future user, then which of the cloud "layers" would you most likely use?

□ Security services in the cloud

□ Individual software packages (SaaS)

□ Infrastructure as a service (laas - storage, network etc.)

□ Platform as a Service (PaaS - complete OS and software packages available via cloud services)

□ Other (please specify)

10. Please say in your own words what particular applications you outsource to the cloud (if you are a user).

11. Please say in your own words what you think are the main vulnerabilities in cloud services.

12. Please state what kind of risk assessment, if any and if known to you, that your organization uses either as user or provider.

Your perceptions of immediate risk

This section asks about **direct and immediate risks** to a system being provided as a cloud service. Section 4 will ask you about longer-term, indirect risks.

This following list outlines **general categories of risk** that will be used as the basis for questions within this section:

- **Denial of cloud service**: for example, when malicious attackers can attack the cloud provider through vulnerabilities in the cloud technology, rendering all its clients unable to use their applications, perhaps in critical situations.
- **Denial of service through neighbours**: for example, when a denial of service attack - either commercially motivated or a politically motivated ‘hacktivist’ attack - is mounted on one user, and another user loses service because it happens to be co-located on the same server or uses the same network resources; or when a badly behaved neighbour disrupts a user’s service through maliciously or accidentally monopolising a shared resource.
- **Privilege escalation**: for example, exploitation of the cloud technology to get privileged access to user’s confidential data or to modify important business rules.
- **Social engineering vulnerabilities**: for example, staff employed by the cloud provider being enticed into disclosing your data to malign third parties, or introducing viruses to your applications.
- **Service migration**: for example, your data being physically located in parts of the world where such data become vulnerable to loss of privacy or to being put in quarantine.
- **User redirection**: for example, individuals trying to get access to your applications are maliciously or accidentally directed elsewhere due to vulnerabilities in the cloud service, and disclose confidential information or become unable to use the service in critical situations.

13. For each of the following general categories of risk, please rate the likelihood that a risk could materialize and the potential consequences if it did.

	Unsure	Very low	Low	Moderate	High	Very high
Likelihood of denial of cloud service	○	○	○	○	○	○
Consequences of denial of cloud service	○	○	○	○	○	○
Likelihood of denial of service through neighbours	○	○	○	○	○	○
Consequences of denial of service through neighbours	○	○	○	○	○	○
Likelihood of privilege escalation	○	○	○	○	○	○
Consequences of privilege escalation	○	○	○	○	○	○
Likelihood of social engineering vulnerabilities	○	○	○	○	○	○
Consequences of social engineering vulnerabilities	○	○	○	○	○	○
Likelihood of service migration	○	○	○	○	○	○
Consequences of service migration	○	○	○	○	○	○
Likelihood of user redirection	○	○	○	○	○	○
Consequences of user redirection	○	○	○	○	○	○

14. For each of the same general categories of risk, please indicate which parties should carry the most responsibility for managing this risk. Multiple answers can be selected.

	Individuals within organizations using cloud services	System managers within organizations using cloud services	Cloud service providers	Third party software developers	Network service providers	Other
Denial of cloud service	□	□	□	□	□	□
Denial of service through neighbours	□	□	□	□	□	□
Privilege escalation	□	□	□	□	□	□
Social engineering vulnerabilities	□	□	□	□	□	□
Service migration	□	□	□	□	□	□
User redirection	□	□	□	□	□	□

Other (please specify)

15. For each of the following system elements, please rate their vulnerability to some kind of attack or failure and please rate the seriousness of the impact this would have

	Unsure	Very low	Low	Moderate	High	Very high
Vulnerability of physical servers and their operating systems	○	○	○	○	○	○
Impact on physical servers and their operating systems	○	○	○	○	○	○
Vulnerability of virtualisation software	○	○	○	○	○	○
Impact on virtualisation software	○	○	○	○	○	○
Vulnerability of application software	○	○	○	○	○	○
Impact on application software	○	○	○	○	○	○
Vulnerability of network services	○	○	○	○	○	○
Impact on network services	○	○	○	○	○	○
Vulnerability of a users' front-end computers	○	○	○	○	○	○
Impact on a users' front-end computers	○	○	○	○	○	○
Vulnerability of a cloud provider organization	○	○	○	○	○	○
Impact on the cloud provider organization	○	○	○	○	○	○
Vulnerability of a user organization	○	○	○	○	○	○
Impact on the user organization	○	○	○	○	○	○
Vulnerability of individuals within the user organization	○	○	○	○	○	○
Impact on the individuals within the user organization	○	○	○	○	○	○

Perceptions of Risk in Cloud Computing for Critical Infrastructures: A

Your perception of less direct risk

This section asks questions about your views on long-term risks in cloud computing.

This following list outlines the **general categories of LONG-TERM risk** that will be used as the basis for questions within this section:

- **Loss of in-house technical capability**: for example, losing all in-house expertise in providing computing services, so that the user organization can no longer cope if there is loss of cloud service for any reason, and loses its ability to assess new computing technologies
- **Loss of data control**: The inability to regain control of data (personal or corporate) held in cloud services.
- **Inability to trace security breaches**
- **Loss of technical choice and direction**: for example, finding that the cloud provider starts to use a technology or adopts a new policy (or extends an existing one) that allows it to mine your data for its own purposes.
- **Resource saturation**: The potential long term loss of performance as cloud computing becomes commonplace. For example, congestion in network traffic and congestion in cloud server capacity.
- **Resource allocation and price escalation**: Risk resulting from the loss of control over how much computing resource the individuals within a user organization is paying for, and the future prices it will be charged
- **Loss of control through cloud provider supply chains**: Loss of control over the computing supply chain through cloud providers themselves relying on other providers for specific computing or network resources
- **Cloud provider unscalability during peak demand**: The inability of cloud provider to cope with concentrated peaks of computing demand. For example during adverse weather conditions.

16. For each of the following general categories of risk, please rate the LIKELIHOOD that a risk could materialize, and the potential CONSEQUENCES.

	Unsure	Very low	Low	Moderate	High	Very high
Likelihood of loss of in-house technical capability	○	○	○	○	○	○
Consequences from the loss of in-house technical capability	○	○	○	○	○	○
Likelihood of loss of data control	○	○	○	○	○	○
Consequences from the loss of data control	○	○	○	○	○	○
Likelihood of the inability to trace security breaches	○	○	○	○	○	○
Consequences from the inability to trace security breaches	○	○	○	○	○	○
Likelihood of loss of technical choice and direction	○	○	○	○	○	○
Consequences from the loss of technical choice and direction	○	○	○	○	○	○
Likelihood of resource saturation	○	○	○	○	○	○
Consequences from resource saturation	○	○	○	○	○	○
Likelihood of resource allocation and price escalation issues	○	○	○	○	○	○
Consequences from resource allocation and price escalation issues	○	○	○	○	○	○
Likelihood of loss of control through cloud provider supply chains	○	○	○	○	○	○
Consequences from the loss of control through cloud provider supply chains	○	○	○	○	○	○
Likelihood of cloud provider unscalability during peak demand	○	○	○	○	○	○
Consequences from cloud provider unscalability during peak demand	○	○	○	○	○	○

17. Please describe any long-term risks you believe exist in cloud computing that are not covered in the above question.

C OpenFlow: A Security Analysis

OpenFlow networks [MAB$^+$08] are being increasingly used to realise the networks of large-scale data centres that support cloud computing services. In this paper, we present a security analysis of OpenFlow, which makes use of a combination of the Microsoft STRIDE vulnerability analysis approach and attack trees. Understanding the security issues associated with OpenFlow networks (and software-defined networks, in general) will become increasingly important as this technology becomes more readily applied to data centres. Furthermore, this paper presents a concrete example of how Data Flow Diagrams (DFDs) can be applied to model a given architecture, such as a cloud computing deployment, which we advocate as part of the risk assessment process that is described in Section 7 of this deliverable.

R. Klöti, V. Kotronis, P. Smith, “OpenFlow: A Security Analysis,” in 8th Workshop on Secure Network Protocols (NPSec 2013), Göttingen, Germany, October, 2013.

A more in-depth description of this work can be found in [Row13].

OpenFlow: A Security Analysis

Rowan Klöti
ETH Zurich
Zurich, Switzerland
Email: rkloeti@ee.ethz.ch

Vasileios Kotronis
ETH Zurich
Zurich, Switzerland
Email: vkotroni@tik.ee.ethz.ch

Paul Smith
AIT Austrian Institute of Technology
2444 Seibersdorf, Austria
Email: paul.smith@ait.ac.at

Abstract—**Software Defined Networking (SDN) has been proposed as a drastic shift in the networking paradigm, by decoupling network control from the data plane and making the switching infrastructure truly programmable. The key enabler of SDN, OpenFlow, has seen widespread deployment on production networks and its adoption is constantly increasing. Although openness and programmability are primary features of OpenFlow, security is of core importance for real-world deployment. In this work, we perform a security analysis of OpenFlow using STRIDE and attack tree modeling methods, and we evaluate our approach on an emulated network testbed. The evaluation assumes an attacker model with access to the network data plane. Finally, we propose appropriate counter-measures that can potentially mitigate the security issues associated with OpenFlow networks. Our analysis and evaluation approach are not exhaustive, but are intended to be adaptable and extensible to new versions and deployment contexts of OpenFlow.**

I. Introduction

Software Defined Networking (SDN) is the key outcome of extensive research efforts over the last decade towards the transformation of the Internet to a more open, programmable, reliable, secure and manageable infrastructure. The main concepts of SDN are: *i)* the separation of the network control plane from the data plane and, *ii)* a logically centralized controller [1], communicating with the data plane over open and standardized interfaces and protocols. The control applications running on top of element *(ii)* see a network-wide view based on the abstraction of the distributed network state.

OpenFlow [2] is a standardized [3] protocol which implements the aforementioned notion of SDN. It is used for the interaction between a network switch, constituting the data plane, and a controller, constituting the control plane. The switch performs packet forwarding using one or more flow tables. These tables contain sets of rules matching to flows traversing the switch (i.e., matching to packet header patterns), corresponding actions (e.g., forwarding or header rewriting), and counters used for measurements. The flow rules are installed on the switch by the controller. The controller can choose to install them *proactively* on its own accord, or *reactively* in response to a notification by the switch regarding a packet failing to match existing rules.

Despite having started as a largely academic endeavour, OpenFlow has been increasingly deployed in production systems over the past two years. For instance, Google has deployed OpenFlow within its datacenter backbone network to maximize utilization on links carrying huge elastic loads [4]. Major vendors such as Cisco, Juniper and HP are offering OpenFlow support in their products [5], and they are using OpenFlow capabilities to differentiate within the growing SDN market [6]. It seems very likely that the adoption of OpenFlow will continue at an increasing rate in the coming years, as service providers and cloud hosts hope to accelerate service deployment, enable easier cloud management and build novel applications on top of their networks [7].

978-1-4799-1270-4/13/$31.00 © 2013 IEEE

Given the potential of SDN in general (and OpenFlow in particular) to revolutionize the way in which networks are managed, looking into the security implications of OpenFlow-based setups while the technology is still young constitutes a very important and challenging task. Although there are research publications on the deployment of security applications *over* OpenFlow [8], [9], none of these address the core issue of the security of the protocol *itself*. To the best of our knowledge, there is no *official* security analysis of OpenFlow available to the public. For the sake of completeness, we note the work in progress described with Internet Drafts [10], which complement our current work. In this paper, we make the following contributions:

a) ***Security Analysis:*** We perform a high-level, extensible and adaptable security analysis of OpenFlow (protocol and network setups), using the STRIDE [11] vulnerability modeling technique. By combining STRIDE with attack tree approaches [12], we provide a fitting methodology for analyzing OpenFlow from a security perspective, uncovering potential vulnerabilities and describing exploits.

b) ***Evaluation:*** We experimentally demonstrate prominent vulnerabilities which are yielded by our security analysis. Further, we implement test-suites in order to exhibit the impact of the exploitation of these vulnerabilities on a widely used OpenFlow virtual switch [13] and controller [14], using an OpenFlow network emulator [15].

c) ***Recommendations:*** Based on our security analysis and evaluation of OpenFlow, we propose techniques that could prevent or mitigate the identified security issues, depending on the deployment and operation context.

The rest of the paper is structured as follows: Section II provides an overview of our security analysis of OpenFlow, along with the methodology used and vulnerabilities found. Section III describes the evaluation environment and presents the results of our test-suite for different attacks. In Section IV we recommend prevention and mitigation techniques stemming from our analysis and evaluation. Section V gives an overview of the related work. Finally, we conclude.

II. OpenFlow security analysis

We have carried out a structured security analysis of the OpenFlow protocol. Here, we provide an overview of the methodology applied to conduct this analysis. For more details we refer the reader to our work in [16], where the full methodology and results are presented.

A. Methodology

To implement the security analysis of OpenFlow, we combine two modeling techniques: Microsoft's STRIDE methodology [11] and attack trees [17]. In an initial phase, the STRIDE methodology is used to construct a model of an OpenFlow system and enumerate its potential vulnerabilities; subsequently, attack trees are employed to explore how the identified vulnerabilities could be exploited by an attacker.

Using STRIDE, a Data Flow Diagram (DFD) of a target system can be developed. This DFD shows the system's components, including processes, data stores, (conditional) data flows and trust boundaries. With a DFD in place an analyst then examines the potential vulnerabilities of each component using the STRIDE mnemonic: **S**poofing, **T**ampering, **R**epudiation, **I**nformation Disclosure, **D**enial of Service, and **E**levation of Privilege. For instance, one might consider the possibility of a Denial of Service (DoS) to an OpenFlow controller process, and evaluate its impact on the overall system. The result of this analysis is a set of system component and vulnerability pairs.

We use attack trees to explore how an identified vulnerability could be exploited. The root of an attack tree is an attacker's ultimate objective – in our case, an OpenFlow component and vulnerability pair, derived by STRIDE. Sub-nodes in a tree represent intermediate attack objectives; leaf nodes represent basic actions and events. Branches in a tree can have logical OR or AND semantics, whereby any sub-node or all sub-nodes must be satisfied to achieve a goal, respectively. The analysis begins at the root node; child nodes are created recursively by decomposing the parent objective.

We made the following assumptions about the attacker's capabilities: they are unable to gain access to the secure control channel that provides connectivity between an OpenFlow switch and its controller, and they cannot directly compromise the system on which the controller or the switch runs. We made these assumptions for two reasons: *(i)* we assume that a network operator has taken reasonable precautions to secure the controller and associated communication channel (e.g., via TLS), and *(ii)* we wanted to focus on threats that emerge from the data plane as a consequence of using OpenFlow.

B. Modeling and Analyzing OpenFlow via STRIDE

Fig. 1 presents a simplified version of a DFD of an OpenFlow switch (for space reasons, we show only a simplified DFD). A number of processes are shown in Fig. 1 that perform forwarding tasks (i.e., *Data path*, implemented on the hardware of the switch) and OpenFlow-related activities: the *OpenFlow Module*, which runs as a software on the switch's CPU and performs tasks such as managing the *Flow table* based on interactions with the controller, and the *Secure Channel* process that handles switch–controller communication. Data flows are defined, e.g., *Read flow table* and *Packet sample*. A trust boundary exists between the data path and the OpenFlow components, as indicated by the dashed-line. Interactions across such boundaries should be carefully considered, as they are likely sources of attacks. Finally, the *Flow table* data store is shown, which contains flow rules for matching L2 – 4 headers, actions to be invoked on flows, and counters.

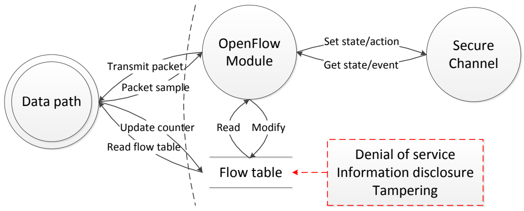

Fig. 1. Simplified DFD for an OpenFlow switch, showing relevant vulnerabilities

With a DFD in place, one can analyze each component using the STRIDE mnemonic. We observe that *Information Disclosure*, *Denial of Service* and *Tampering* vulnerabilities and attacks are possible. An attack with severe consequences is a *Denial of Service* against the flow table, whereby an attacker aims to overload the table with flow rules, illustrated in Fig. 1. We show how an attacker can achieve this in Sec. III-B1. We further note the possible detrimental effects of such attacks on the controller as well as the secure channel – in the case of the former, the attack may also affect further switches managed by the same controller. If the attacker has sufficient knowledge of the internal implementation, they may be able to effect a hash collision attack on the flow table or analogous data structures in the controller. With respect to *Information Disclosure*, we note that by observing differences in controller response times, an attacker may be able to derive information about network state, such as active flow rules. Sec. III-B2 gives an example of such an attack. Furthermore, with respect to *Tampering* we mention the possibility of cache poisoning attacks against the flow table and/or controller state. More complex attacks that combine the aforementioned primitives (e.g., using knowledge acquired by a preemptive Information Disclosure attack in order to mount an effective DoS) can also be formulated.

C. Attack Tree Analysis

We have developed attack trees for several vulnerabilities that were identified using the STRIDE methodology. An example attack (sub-)tree is shown in Fig. 2, which shows how an attacker might implement an Information Disclosure attack against an OpenFlow controller. In this scenario, an attacker is attempting to learn the nature of the controller's behavior, e.g., whether and which aggregated rules are in use for certain flows, by measuring the time it takes for selected packets to reach an end-point and return. The intuition for this attack is that packets which do not correspond to already installed flow rules require forwarding to the controller, thus inducing an additional forwarding and processing delay.

Fig. 2 shows the steps necessary to realize this attack – an attacker must elicit a response from an end-point, either by gaining access to multiple clients (e.g., by compromising a machine) or forcing a client to reproduce a response. Either of these options is possible, as indicated by the logical OR

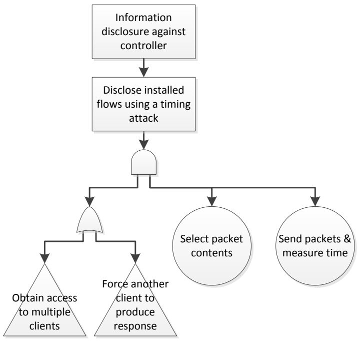
Information disclosure against controller
Disclose installed flows using a timing attack
Obtain access to multiple clients
Force another client to produce response
Select packet contents
Send packets & measure time

Fig. 2. Simplified attack (sub-)tree showing an Information Disclosure attack against an OpenFlow controller

branch in the attack tree. Subsequently, an attacker selects the packet contents associated with the information they wish to disclose, sends the packet and measures the round-trip time. A more detailed description of this attack is given in Sec. III-B2, wherein we describe an experimental implementation of it.

III. Empirical evaluation

A. Setup and Emulation Environment

In this section, we provide an overview of the emulation environment, the traffic generation tools and the network setup that we used in order to evaluate the Denial of Service and Information Disclosure vulnerabilities and consequent attacks. Further evaluation details and scripts implementing our test-suite are provided in [16].

1) Emulation Environment: We used the *Mininet* framework to create virtual networks based on Open vSwitch [13]. Mininet utilizes network namespaces, a feature of the Linux kernel, to implement lightweight network virtualization. Individual clients are modeled as *nodes* (which can be *hosts*, *switches* or *controllers*) and possess *interfaces*, representing NICs. Virtual links between *interfaces* are modeled as *links*, which may be subject to performance constraints, such as bandwidth, delay, buffer size and simulated packet loss. See [18] for more information on the Mininet implementation.

2) Traffic Generation: To implement the attacker, the packet generation and analysis framework *scapy* is used. It is a Python-based framework allowing the creation of packets with arbitrary data in the header fields. The utilitity *netcat* is used to emulate a TCP client and server.

3) Network Setup: The main setup consists of two identical client systems, a user-space OpenFlow switch and a POX-based controller [14]. This setup is depicted in Fig. 3. Each node has a unique virtual network connection to the switch. The attacker controls one or more client systems. The attacker does not have any control over or access to the switch or the controller. External observations (e.g., packet dumps between the switch and the controller) are not permitted for the attacker, but may be used to evaluate the impact of the attacks. Some forms of attack require a more sophisticated network environment depicted in Fig. 4, which shows two virtual switches linked together. Each of the switches is connected to three further virtual hosts. A single controller controls both switches. As above, all of the data path links have identical performance parameters, while the control path links also have identical and distinctive (from the data path links) performance characteristics. This setup requires that the controller supports layer-3 forwarding properly.

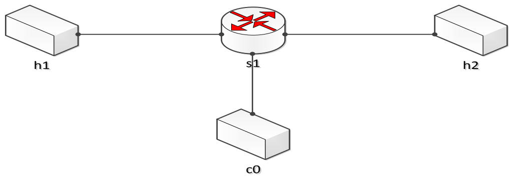

Fig. 3. Schematic diagram of virtual network setup used in Sec. III-B1

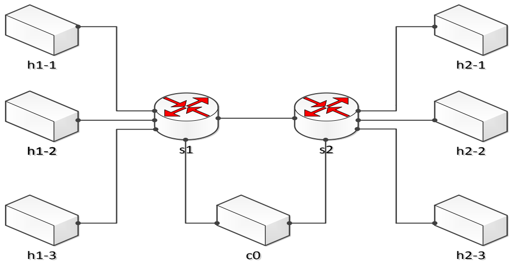

Fig. 4. Schematic diagram of virtual network setup used in Sec. III-B2

B. Results

1) Denial of Service: The objective of this attack is to generate a large number of packets that will be sent to the controller and result in it installing a new flow rule for each packet, eventually overflowing the flow table. We utilize the POX module *forwarding.l2_learning* which implements a layer-2 learning switch, exemplifying a purely reactive strategy. As the controller only installs rules matching header fields *exactly*, it is only necessary to permute some value in a packet header to cause the installation of a new flow rule. For this purpose, the source and destination port fields of UDP packets are used. The *forwarding.l2_learning* module has been modified to accept user-provided soft timeout values, which determine when a flow rule expires, so that their effect on the attack may be observed. The effect of the attack is measured by packet loss and instances of the *All tables full* error being produced by the switch. These correlate perfectly, so only the former is shown here. Fig. 5 shows a steady increase in the number of lost packets with an increasing timeout value. This increase can be explained as follows: larger timeouts mean more persistent flow rules within the table and larger probability of table overflows and denied rule installations. There is also a long plateau between approximately 37 *s* and 67 *s*, probably an artifact of the Open vSwitch implementation [13].

Fig. 6 illustrates that lower performance on the control link tends to aggravate the effect of the aforementioned DoS attack, with the plateau of packet loss being reached earlier (with 31 *s* timeout). The packet loss also exceeds this plateau for lower timeout values (about 64 *s*). There is significantly higher incidence of packet loss being observed for smaller timeout values than in the previous case. This is counterintuitive: we expected that with the same timeouts, a slower control link would result in less flow rules being installed per time unit, thus reducing the incidence of table overflows.

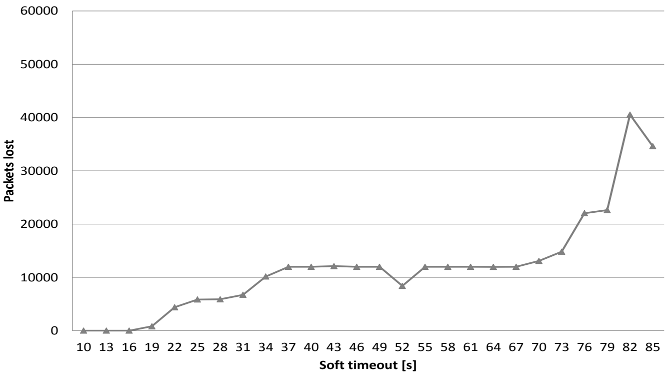

Fig. 5. Test with data link at 100 Mbps, 10 ms delay, control link at 100 Mbps, 1 ms delay

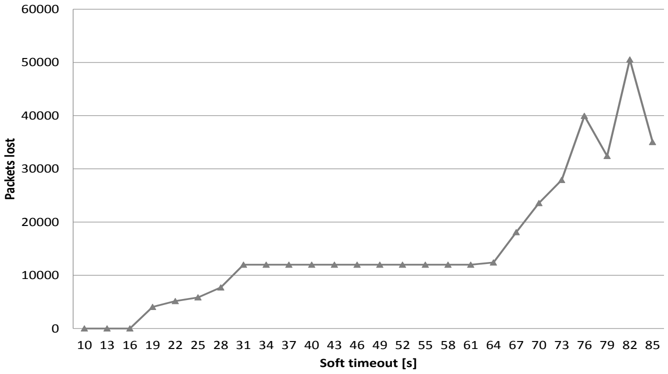

Fig. 6. Test with data link at 100 Mbps, 10 ms delay, control link at 10 Mbps, 10 ms delay

2) Information Disclosure: The objective of the attack is to exploit the use of flow aggregation in order to discover some aspect of network state that would otherwise not be visible to an attacker. This information could be used by an attacker to determine the presence and nature of services on a network. Such knowledge might also be used in a later stage of an attack. The network setup used here is described in Fig. 4. If a server is connected to the second switch (*s2*), and several clients to the first switch (*s1*), then the aggregation occurring in *s1* in response to several connections from the clients to the server could allow another client connected to *s1* to deduce that such a connection exists. This is performed by timing the TCP setup; if a second connection attempt is substantially faster than the first, then a new flow rule was installed in response to the connection attempt. Conversely, if there is no significant difference, the attacker may conclude that a flow rule already existed. For this attack to be performed, it is necessary that dynamic aggregation of flow rules is in use. This is achieved with the POX module *forwarding.l3_aggregator_simple*. This module sets the following header values to wildcards: link and network layer source addresses, transport layer source port and the physical switch port. The forwarding behaviour does not need to depend on source values, so the aggregation of these fields is reasonable to minimize the number of flow rules.

We measure the distribution of setup times in order to determine the certainty with which we may conclude whether an existing flow rule is present. We also perform the operation with a non-aggregating controller, acting as the control. This allows us to observe how significant the differences in timing are. Fig. 7 shows a histogram of the control data. The two data sets exhibit the case before a parallel connection is created from another client to the server, and the case afterwards. In the latter case, aggregation may occur, if the controller allows this. In the control, this is *not* allowed. The distributions here are equal within a reasonable tolerance, as expected.

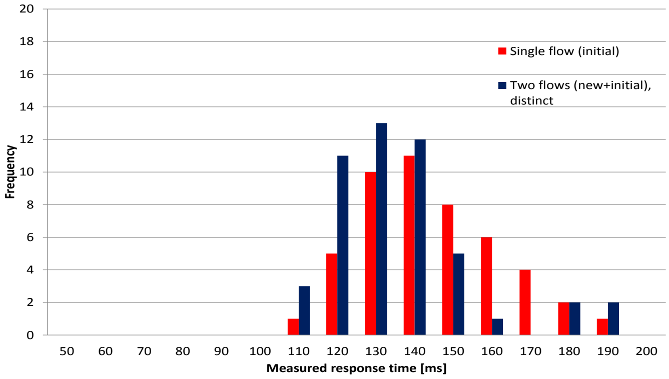

Fig. 7. Histogram of control using *forwarding.l3_learning* controller with symmetric timing at 10 ms

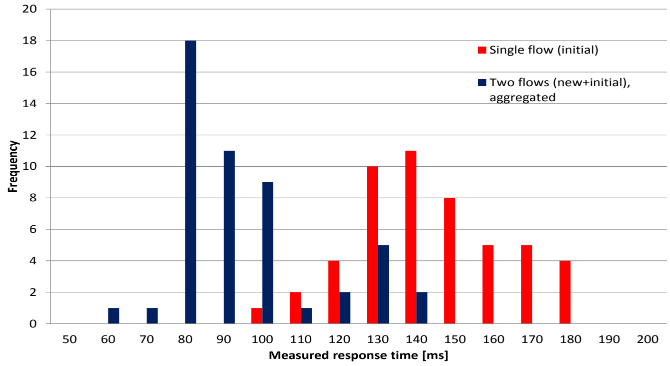

Fig. 8. Histogram of data using *forwarding.l3_aggregator_simple* controller with symmetric timing at 10 ms

Fig. 8 shows a histogram of measured times when aggregation is in effect, for a network with symmetric delays (the latency is the same on all control and data links). The distribution of the second data set (with aggregation) is clearly distinguishable from the pre-aggregated one, in contrast to the previous case. This attack is dependent on the latencies of the network in question; longer latencies on the control path increase the distinguishability of the distributions, while longer latencies on the data path or multiple hops diminish it.

IV. Recommendations

Based on the findings of our model-based analysis in Sec. II and experimental results in Sec. III, we provide insights regarding techniques that could potentially counter the uncovered security issues within OpenFlow deployments.

To organize our recommendations, we consider the various network setups in which OpenFlow may be deployed, together with their special characteristics such as: 1) the user base (e.g., known and trusted or external and untrusted), 2) the direction of flow establishment (e.g., inbound from untrusted sources or outbound from trusted insiders), and 3) the operational requirements: *security* (e.g., prevention of unauthorized external access), *performance* (e.g., throughput and latency) and *reliability* (e.g., minimization of downtime or fast fail-over capability). Table I includes usual network types together with their corresponding properties and requirements. Requirements are ranked with high (H), medium (M) or low (L) importance.

Type	User base known	Flows established (main direction)	Requirements Security	Performance	Reliability
Corporate	✓	Outbound	**H**	**M**	**H**
Academic	✓/✗	Outbound	**M**	**L**	**M**
Research	✓	Both	**L**	**L**	**L**
Data center	✓/✗	Both	**H**	**H**	**H**
Backbone	✗	Both	**L**	**H**	**H**
DMZ	✗	Inbound	**H**	**M**	**H**
Special purpose	✓/✗	Unknown	**M**	**M**	**M**

TABLE I. DIFFERENT NETWORK TYPES AND THEIR PROPERTIES.

Next, we mention proposed state-of-the-art applications [7] of OpenFlow within the aforementioned network environments: 1) dynamic or proactive switching and routing, 2) multicasting, 3) access control, 4) load balancing, 5) fail-over and path recovery, 6) QoS policy enforcement, 7) network virtualization and isolation [19], and 8) monitoring and instrumentation. Different network types may benefit more from certain applications, e.g., a monitoring controller application that captures the behavior of new flows is valuable to datacenter and DMZ networks, as well as research networks. We note that the prevention and mitigation techniques to be used depend on the combination of the network type and application: there is no "one-size-fits-all' recipe or practice.

In the following two sections, we describe a number of useful techniques that can potentially mitigate DoS and Information Disclosure attacks, applicable to diverse networks. We note that the context of these approaches is not a generic network setup, but an *OpenFlow environment*. Although these techniques are *recommended*, they require further investigation and empirical evaluation, beyond the scope of this paper.

A. Denial of Service

DoS attacks can target the controller and/or the switch, aiming at crippling the communication between the components or the components themselves. Mitigation in this context is coupled with the continuing operation of those elements, without noticeable performance degradation. The following approaches are conceivable for achieving this goal.

1) Rate Limiting, Event Filtering, Packet Dropping and Timeout Adjustment: Rate limiting on the control channel and/or the data interface can allow the controller and/or the switch respectively to remain responsive during a DoS attack, although it cannot protect other users from negative effects. Event filtering enables the selective handling of event types by the controller, potentially increasing system resilience. Furthermore, in case the attacker can be detected with sufficient precision, flow rules that match the malicious traffic can be installed on the switch, effectively dropping the misbehaving packets. Even if the administrators are unable to isolate the attacker, traffic prioritization and QoS mechanisms can be put in place to cope with the load. Lastly, flow timeouts can be tuned to decrease the impact of DoS accordingly, since larger timeouts can lighten the load on the control channel while shorter ones can decrease the number of switch flow table overflow incidents. Some of these approaches have been standardized by the ONF [3] in recent versions of OpenFlow.

2) Flow Aggregation: Flow Aggregation is a proactive strategy where each flow rule matches multiple network flows, thus reducing the number of rules required to match network traffic. Its advantage is that the flow table is less prone to overflows, while the controller receives less load on the control channel, i.e., fewer unmatched packet notifications by the switch. Of course, this comes with the cost of precision and responsiveness. Aggregated flow rules are suitable for networks that practice proactive strategies, e.g., backbone carriers, but they may not be applicable to enterprise networks, where fine-grained flow control is a key security objective. This method is obviously more effective when the attack traffic has limited dispersion characteristics.

Proposed measures	Implemented on			Suited for
Description	Switch	Controller	Protocol	Network
Rate limiting	✓	✓	✓	All
Event filtering	✓	✓	✓	
Packet dropping	✓	✗	✗	
Reduce timeouts	✗	✓	✗	
Flow aggregation	✗	✓	✗	Backbone Data center DMZ
Attack detection	✓	✓	✗	Corporate Academic DMZ
Access control	✗	✓	✗	Corporate Special cases

TABLE II. PROPOSED COUNTERMEASURES AGAINST DENIAL OF SERVICE ATTACKS AND THEIR CONTEXT.

3) Attack Detection: Detecting DoS is itself a very difficult problem and open research area [20], [21]. Here, we note that basic detection functionality could be implemented as a logically centralized controller application. Related performance issues (such as control channel latency) can be technically dealt with in part via physical distribution or multi-thread processing. On the switch's side, the flexible forwarding behavior of OpenFlow could be employed to direct flows to monitored paths for processing, while the monitoring systems themselves can "subscribe", via OpenFlow, to the type of traffic they want to examine dynamically. The addition of OXM to OpenFlow v1.2 [3] offers useful extensions to implement DPI on a flow at relatively low performance costs, subject to vendor support.

4) Access Control: Enforcement of access control lists in the form of flow rules on the table of the OpenFlow switch is also a feasible and low-cost approach. For example, traffic originating from inside the trusted domain may be allowed to pass, while inbound traffic would be compared against a whitelist set of flow rules. This solution is worth considering for corporate networks, in which traffic is likely to originate from internal hosts or trusted external ones (e.g., over VPN connections). On the other hand, it is not bound to be applicable in DMZ or backbone networks. Lastly, directing flows to actual firewalls and IPS that analyze and filter traffic is another solution, although these systems do not separate the data and control planes and do not follow the SDN principles of OpenFlow. The problem of detecting or predicting malicious traffic still remains. In any case, the controller is responsible for installing the appropriate flow rules that handle such traffic (e.g., which drop it), proactively or reactively on the switch.

Table II summarizes the diverse approaches that can be employed to mitigate DoS attacks, along with the appropriate implementation context (switch, controller and protocol) and the applicability on different network environments.

B. Information Disclosure

Information Disclosure, arising from timing analysis, can reveal certain aspects of a network's state as well as a controller's strategy to an attacker. Mitigation in this context means ensuring that the observable system parameters do not expose the internal system state. For example, the increased delay for the establishment of a new flow rule in response to an incoming packet can inform the attacker about the behavior of the OpenFlow controller. The following approaches are conceivable for achieving mitigation.

1) Proactive Strategies: Proactive flow rule establishment removes the dependency of the response time on the network state (i.e., the switch flow table entries). Of course, automatic flow aggregation techniques may worsen the situation, since an attacker may infer the presence of another connection that is aggregated with his current one from the switch's perspective.

2) Randomization: Increasing the variance of measurable response times can increase the statistical uncertainty of the attacker and reduce the strength of the attack considerably. A way to implement this via OpenFlow is to randomize the timeouts of the installed flow rules, in order to mimic an unpredictable behavior that will prevent the attacker from forming a coherent view of the network state. In this case, trade-offs between the level of timing obfuscation and performance degradation need to be carefully evaluated.

3) Attack Detection: Any attack that is based on timing analysis is likely to exhibit a distinctive, repetitive pattern that may be used by a controller application to detect it, enact counter-measures or notify an administrator. Counter-measures could include dropping suspicious traffic, introducing randomization or adapting the forwarding strategy accordingly.

V. Related Work

To the best of our knowledge, there is no prior *official* security analysis of OpenFlow itself, [10] notwithstanding. [8] introduces an extension called *FortNOX* to the NOX controller [1], providing a role-based access control system that validates digitally signed flow rules before table insertion. *FortNOX* is focused on the control plane and proposes extensions to OpenFlow control, while our work is focused on the data plane and is an *analysis* of OpenFlow. [19] proposes *FlowVisor*, a system allowing virtual networks to be built on top of an OpenFlow network, thus enabling multiple experimental network slices that do not interfere with production traffic. [9] proposes *VeriFlow*, a system used to validate the forwarding behavior of an OpenFlow network in real time. [22] describes *OpenFlow Random Host Mutation*, a technique that exploits OpenFlow to protect end systems from attacks by providing them with virtual external IP addresses, translated into the actual ones by the controller. [23] describes an application of OpenFlow for the detection of DDoS attacks, making use of *Self Organising Maps* to classify traffic patterns.

VI. Conclusions

We presented a security analysis and modeling methodology for the OpenFlow protocol and network setups. Using STRIDE [11] and data flow diagrams we uncovered vulnerabilities such as Denial of Service and Information Disclosure which are exacerbated due to the nature of SDN. These vulnerabilities were developed into feasible attacks through attack tree modeling methods. The feasibility and impact of the attacks were evaluated using network emulation, testing tools, an open-source controller and a virtual OpenFlow switch distribution. Based on our analysis and evaluation, we recommended numerous prevention and mitigation techniques corresponding to different network deployment and operation contexts. Our methodology and testing approach can be adapted to future versions and extensions of OpenFlow. We hope that this work will help SDN researchers [24] and the OpenFlow standardization body [3] in the ongoing effort [10] for SDN architectures, applications and standards that are more secure *by design*.

Acknowledgments

This work has been partially funded by the EU FP7 projects OFELIA and SECCRIT.

References

[1] NOXRepo.org, "NOX," http://www.noxrepo.org/.

[2] N. McKeown *et al.*, "OpenFlow: enabling innovation in campus networks," *SIGCOMM CCR*, Mar. 2008.

[3] "Open Networking Foundation," https://www.opennetworking.org/.

[4] U. Hoelzle, "OpenFlow@Google," http://opennetsummit.org/archives/apr12/hoelzle-tue-openflow.pdf.

[5] "Open Networking Summit," http://www.opennetsummit.org/.

[6] "SDNCentral Exclusive: SDN Market Expected to Reach \$35B by 2018," http://www.sdncentral.com/sdn-blog/sdn-market-sizing/2013/04/.

[7] "SDN applications," http://searchsdn.techtarget.com/resources/SDN-applications.

[8] P. Porras *et al.*, "A security enforcement kernel for OpenFlow networks," in *Proc. of HotSDN*, 2012.

[9] A. Khurshid *et al.*, "VeriFlow: verifying network-wide invariants in real time," in *Proc. of HotSDN*, 2012.

[10] "Security Analysis of the ONF OpenFlow Switch Specification," http://tools.ietf.org/html/draft-mrw-sdnsec-openflow-analysis-02.

[11] S. Hernan *et al.*, "Uncover Security Design Flaws Using The STRIDE Approach," http://msdn.microsoft.com/en-gb/magazine/cc163519.aspx, 2006.

[12] V. Saini *et al.*, "Threat modeling using attack trees," *J. Comput. Sci. Coll.*, Apr. 2008.

[13] "Open vSwitch," http://openvswitch.org/.

[14] NOXRepo.org, "About POX," http://www.noxrepo.org/pox/about-pox/.

[15] "Mininet," http://mininet.github.com/.

[16] Rowan Klöti, "OpenFlow: A Security Analysis. MSc thesis, D-ITET, ETH Zurich," ftp://ftp.tik.ee.ethz.ch/pub/students/2012-HS/MA-2012-20.pdf, 2013.

[17] P. Khand, "System level security modeling using attack trees," in *Proc. of the 2nd Intl. Conf. on Computer, Control and Communication*, 2009.

[18] N. Handigol *et al.*, "Reproducible network experiments using container-based emulation," in *Proc. of ACM CoNEXT*, 2012.

[19] R. Sherwood *et al.*, "Carving research slices out of your production networks with OpenFlow," *SIGCOMM CCR*, Jan. 2010.

[20] G. Thatte *et al.*, "Parametric methods for anomaly detection in aggregate traffic," *IEEE/ACM Trans. Netw.*, Apr. 2011.

[21] J. Cheng *et al.*, "DDoS attack detection method based on linear prediction model," in *Proc. of ISIC*, 2009.

[22] J. H. Jafarian *et al.*, "Openflow random host mutation: transparent moving target defense using software defined networking," in *Proc. of HotSDN*, 2012.

[23] R. Braga *et al.*, "Lightweight DDoS flooding attack detection using NOX/OpenFlow," in *Proc. of IEEE LCN*, 2010.

[24] "OpenFlowSec.org," http://www.openflowsec.org/.

D Anomaly Detection in the Cloud

One of the categories of threat that we have identified relates to the challenge of implementing existing security controls in the cloud environment (see Section 6.1.10). An increasingly important security control is anomaly detection systems, which aim to identify deviations from normal behaviour that could indicate an attack or some other challenge, such as a misconfiguration. As a demonstration of this issue, in this paper, we explore the impact that wide-area virtual service migration (which can be applied to increase the fault-tolerance of cloud-based services) has on contemporary network flow-based anomaly detection techniques. The paper suggests that this form of migration can affect these techniques, e.g., resulting in higher numbers of false alarms, pointing the way for further exploration of this problem. In the SECCRIT project, this issue is addressed more comprehensively in the context of work package four, and presented in Deliverable D4.1 [SuhSHB13].

K. Adamova, D. Schatzmann, B. Plattner, P. Smith, "Network Anomaly Detection in the Cloud: The Challenges of Virtual Service Migration," submitted to IEEE International Conference on Communications (ICC), Sydney, Australia, June, 2014.

A more in-depth description of this work can be found in [Ada13].

Network Anomaly Detection in the Cloud: The Challenges of Virtual Service Migration

Kirila Adamova, Dominik Schatzmann, and Bernhard Plattner
ETH Zürich
Zürich, Switzerland
Email: kirila.adamova@gmail.com
{dominik.schatzmann, plattner}@tik.ee.ethz.ch

Paul Smith
AIT Austrian Institute of Technology
2444 Seibersdorf, Austria
Email: paul.smith@ait.ac.at

***Abstract*—The use of virtualisation technology in the cloud enables services to migrate within and across geographically diverse data centres, e.g., to enable load balancing and fault tolerance. An important part of securing cloud services is being able to detect anomalous behaviour, caused by attacks, that is evident in network traffic. However, it is not clear whether virtual service migration adversely affects the performance of contemporary network-based anomaly detection approaches. In this paper, we explore this issue, and show that wide-area virtual service migration can adversely affect state of the art approaches to network flow-based anomaly detection techniques, potentially rendering them unusable.**

I. Introduction

Cloud computing has proved to be a popular way for organisations to provision services for their users. There are a number of reasons for this popularity, including potential reductions in operating costs, flexible and on-demand service provisioning, and increased fault-tolerance. Drawn to these benefits, operators of critical information infrastructures – the ICT infrastructures that support gas and electricity utilities and government services, for example – are considering using the cloud to provision their *high assurance* services. This is reflected in a recent white paper produced by the European Network and Information Security Agency (ENISA), which provides specific guidelines in this area [1].

Deploying high assurance services in the cloud increases cyber-security concerns – successful attacks could lead to outages of key services that our society depends on, and disclosure of sensitive personal information. To address these concerns, a range of security measures must be put in place, such as cryptographic storage and network firewalls. An important measure is the ability to detect when a cloud infrastructure, and the services it hosts, is under attack via the network, e.g., from a Distributed Denial of Service (DDoS) attack. A number of approaches to network attack detection exist, based on the detection of *anomalies* in relation to normal network behaviour [2].

One of the essential characteristics of cloud computing is the use of virtualisation technology, which supports the migration of services across a physical infrastructure within and between large-scale cloud data centres – known as *local* and *wide-area* migration, respectively. The reasons for service migration are manifold, including responding to hardware faults, planned maintenance tasks, and handling localised peaks in service requests by moving services “closer” to their user. Whilst virtual service migration has a number of benefits, it has the potential to make the implementation of security measures challenging, therefore introducing new vulnerabilities [3].

In this paper, we are specifically interested in examining the affect virtual service migration has on contemporary network anomaly-based attack detection techniques – as services move, migration may be observable in the network traffic that is being used for anomaly detection. Such techniques aim to detect anomalous traffic in relation to a learned baseline, which represents normal behaviour. It is unclear to what extent virtual service migration, which is arguably representative of “normal” cloud behaviour, can be incorrectly observed as an anomaly, and therefore an attack. Conversely, attacks may be missed because of virtual service migration. If this problem is significant, anomaly detection techniques could be rendered unusable for the cloud, thus representing a significant vulnerability and a potential inhibitor to the deployment of high assurance services.

Using a novel toolchain, which simulates attacks and virtual service migration in network flow traces, we have examined the detection performance of two anomaly detection techniques – Principal Component Analysis (PCA) [4], [5], [6] and the Expectation-Maximisation (EM) clustering algorithm [7], [8]. In previous research, these detection techniques have been shown to give acceptable detection performance results in non-cloud settings. Under different attack and virtual service migration scenarios, we have measured their ability to reliably detect attack behaviour in the cloud. Our results suggest that, in some configurations, a potentially insecure number of attacks are missed, and an unusably high number of alarms pertaining to normal behaviour are generated. This result draws into question the use of these techniques, and potentially others, in large cloud data centres, in which virtual service migration is a common undertaking.

The rest of this paper is organised as follows: Section II discusses related work – our investigations indicate that, to the best of our knowledge, there is no previous work that directly addresses the problem explored in this paper. A discussion on virtual service migration and its affect from a network perspective is presented in Section III. Section IV describes the toolchain and traffic data that we used to obtain the experimental results, which are described in Section V. We conclude and discuss potential solutions to the problem we have explored in Section VI.

II. Related Work

A number of approaches can be used to detecting network-borne attacks, such as Distributed Denial of Service (DDoS) attacks, to large cloud data centres. There has been significant research interest in and deployment of algorithms that aim to identify deviations from normal traffic behaviour – *anomalies* – that are indicative of attack behaviour. A survey of anomaly detection approaches has been produced by Chandola *et al.* [2]. Our research is based on contemporary approaches that use spectral analysis to detect anomalies in entropy measures, derived from network flow summary data. Tellenbach *et al.* examine the use of Kalman filter, Principal Component Analysis (PCA), and Karhunen-Loève Expansion (KLE) when considering the affect of different entropy measures on detection performance [4]. For our investigation, we use PCA, as it has been shown to perform well when configured appropriately [6] and continues to be investigated in the research community [5]. To explore the potential extent of the problem that we have identified, we have used an anomaly detection algorithm that uses clustering; specifically, the Expectation-Maximisation (EM) algorithm, which has shown to give promising detection performance [7], [8].

Our investigation relates to identifying potential shortcomings in flow-based anomaly detection techniques, which manifest due to their deployment; in this case, wide-area virtual service migration in large cloud data centres. This line of enquiry is, in-part, motivated by previous research, conducted by Brauckhoff *et al.*, which examined the impact sampling of network flow data has on anomaly detection [9]. Their study showed that statistical techniques that identify anomalies in traffic volumes perform less effectively under sampling conditions. Furthermore, they suggest spectral-based analysis, using entropy measures of traffic feature distributions, e.g., source and destination IP address and port numbers, are more robust to sampling. As we will discuss in Section III, wide-area virtual service migration manifests as a change in network traffic volume, observable at a data centre – this is similar to the affect sampling has. This observation was one of the motivations for the choice of a PCA-based approach for the study we present in Section V. To the best of our knowledge, our investigation is the first to examine the affect virtual service migration has on network flow-based anomaly detection techniques.

III. Preliminaries – Virtual Service Migration

In order to consider the analysis results that are presented in Section V, a brief discussion on virtual service migration is required, along with its affect on data centre network traffic.

A. Virtual service migration in large cloud data centres

An essential characteristic of cloud computing is the use of virtualisation technology [10], whereby services execute in virtual machines atop of a physical compute, storage and network infrastructure. Virtualisation supports the ability to migrate services between different underlying physical infrastructure. There are multiple reasons to migrate services as a consequence of the *day-to-day* operation of a large data centre, including load balancing, failure of the underlying hardware [11], in response to routine maintenance tasks, and reducing network costs. In some cases, a service must be migrated between geographically and topologically distinct data centres. For instance, if the majority of client connections for a service originate in Asia, but the service is hosted in Europe, network costs can be decreased if the service is moved topologically closer to its clients. Furthermore, geographical diversity of data centres is supported by commercial cloud providers, such as Amazon, in order to improve fault-tolerance. Moving services between data centres is known as *wide-area* migration. Conversely, *local-area* migration occurs when a service is migrated within a data centre. In both cases, there are multiple approaches to ensure the network traffic that is destined for a migrated service is forwarded to the correct location [12].

B. Observing migration in network traffic

Importantly for anomaly detection, service migration may result in observable effects in network traffic – this depends on where network traffic that is to be analysed is collected in the data centre topology, and the type of migration that is carried out. For example, if a local-area migration is executed, the change in traffic could be observable at Top-of-Rack (ToR) and aggregate switches, but not at the gateway to the data centre. (This is the case if we assume the data centre topology outlined in [13].) Since it is common practice to analyse traffic at the edge of a data centre, local area migration is opaque at this location. However, for wide-area migration, traffic destined to the migrated service will be forwarded to a different data centre, and will stop being received at the origin after the migration process has finished. This will result in potentially observable effects in traffic collected at the edge of a data centre, and hamper anomaly detection techniques. In this paper, it is precisely this problem that we explore – i.e., the affect that wide-area service migration has on anomaly detection techniques when network traffic that is used for analysis is collected at the data centre edge.

Broadly speaking, what occurs at a data centre after a wide-area migration can be described as *removing* all the traffic related to a migrated service at the source data centre, and *adding* new traffic for a service or set of services at the target data centre. We acknowledge there will be other observable effects, such as virtual machine state being transferred, but we expect these to be relatively minor in a large data centre. These effects are illustrated in Fig. 1, which shows changing source port entropy values over time, with a dataset that includes a port scan that operates for its duration and a service migration at time $10/24$. These plots were created using the dataset and toolchain we describe in Section IV. The migration can be seen as the entropy decreases at the source data centre (Fig. 1(a)), where a service is removed and the source port distribution is not as previously dispersed. At the target data centre (Fig. 1(b)), the inverse can be observed, including port entropy changes that could result in false alarms. Another interesting observation that can be made is how the injected anomalies appear after migration – in Fig. 1 these are shown as the peaks. At the source data centre, after migration, the anomalies appear more pronounced when compared to the rest of the traffic. The inverse can be observed at the target data centre after migration.

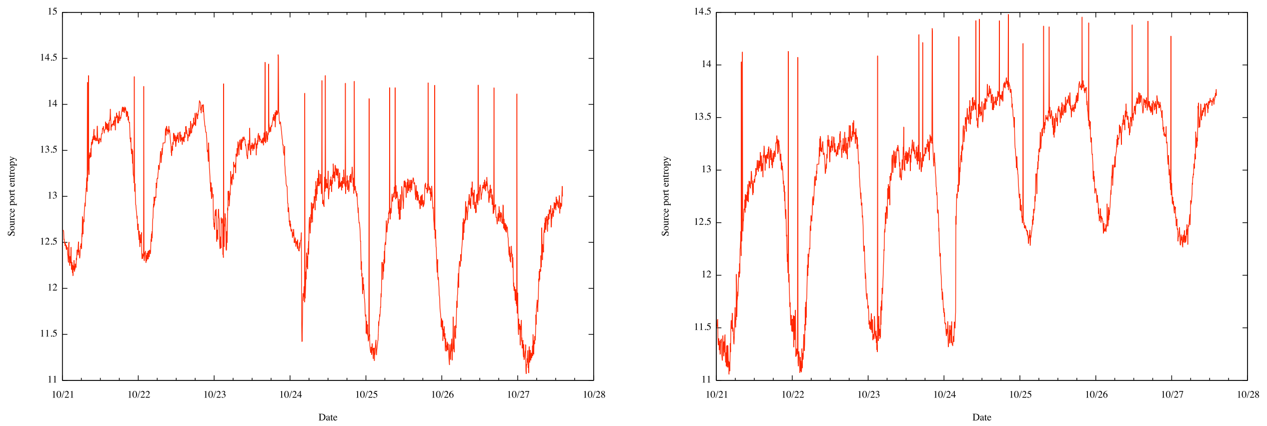

(a) Source data centre

(b) Target data centre

Fig. 1. The affect of migration on source port distribution of traffic at the source and target data centres. Migration occurs at 10/24; the dataset includes a port scan for its duration.

IV. Simulating Cloud Service Migration

In order to examine the affect virtual service migration has on anomaly detection, we created a toolchain that can be used to "simulate" virtual service migration and attacks in a given network flow dataset. The toolchain integrates a number of existing software to achieve this goal. It was necessary to develop the toolchain as we were *(a)* unable to find any public datasets that contained migration behaviour for a large data centre; and *(b)* it afforded us a degree of flexibility for our experiments, e.g., to inject migration at arbitrary times in a baseline dataset, with varying numbers of services being migrated. We briefly describe the dataset we used and our toolchain.

A. *Traffic dataset*

For our experiments, we used network flow data that was collected from the Swiss national research network – SWITCH[1] – a medium-sized backbone network that provides Internet connectivity to several universities, research labs, and governmental institutions. The flow data was collected, without explicit sampling configured, at the border routers of the SWITCH network. We used a week's worth of data from 21–28 October 2012 for our experiments. To produce flow data that is representative of that seen at a large data centre, we extracted the top 200 TCP-based services, based on total number of flows, from the dataset. Services were identified using the techniques described in [14]. We confirmed these services represented those seen at a cloud data centre by comparing their features to a subset which accessed services at one of Amazon's European data centres. This is possible as Amazon publishes the IP address ranges of their cloud data centres. We found them to be comparable. Using this baseline dataset we can then include migration and attack behaviour using our toolchain.

B. *Experimentation toolchain and method*

Fig. 2 presents an overview of the toolchain that we developed for our experiments. Initially, NetFlow data is processed using an extended version of a flow processing framework, called Flowbox[2]. This process includes filtering the top 200 selected services, and storing the flow data for each service in separate files that span five minute periods. Anomalies are generated using the FLAME tool [15], and stored in NetFlow format. We conducted experiments with a volume-based attack, i.e., a Distributed Denial of Service (DDoS) attack, and non-volume based vertical and horizontal port scans. Similarly, anomalies that are to be injected into the dataset are stored in five-minute files.

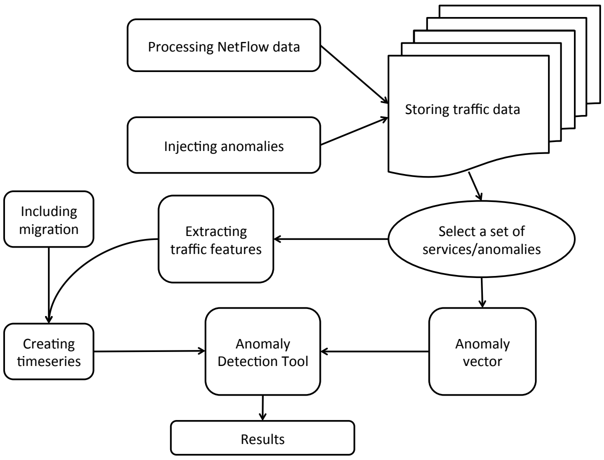

Fig. 2. Overview of the toolchain used for our experiments

The *anomaly detection tool* that analyses the traffic data, along with the injected anomalies and migration, has three inputs: an *anomaly vector*, which specifies at what times anomalies were injected, and two *time series* files – a baseline and evaluation file. A time series file contains entries that summarise each five minute period of the traffic. This description includes seven traffic features and a timestamp. These time series are created by extracting the traffic features from a set of pre-selected services and anomalies. The traffic

[1] http://www.switch.ch/

[2] https://github.com/FlowBox

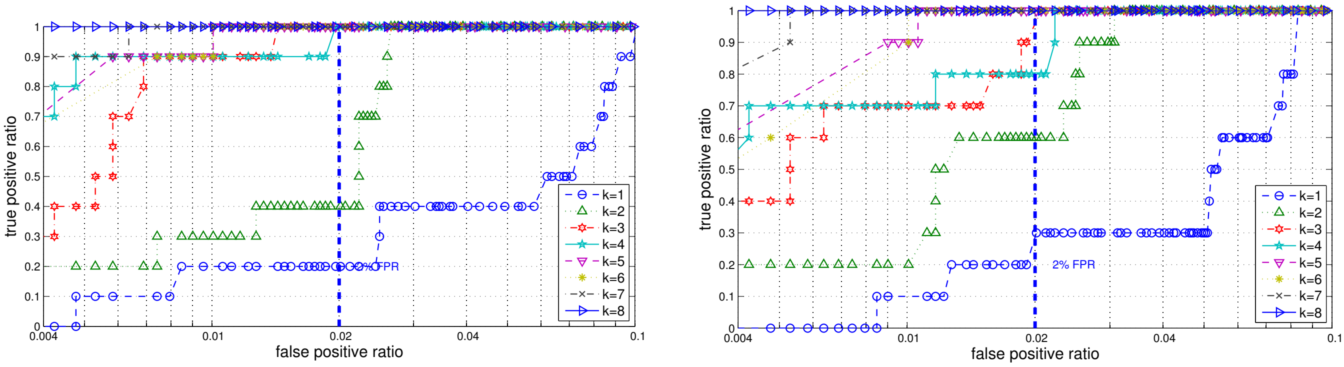

(a) ROC curve *without* virtual service migration

(b) ROC curve *include* virtual service migration

Fig. 3. ROC curves that show the performance of a Principal Component Analysis (PCA)-based anomaly detection approach to detect a Distributed Denial of Service (DDoS) attack, both without (3(a)) and with (3(b)) virtual service migration consisting of 20% of the data centre traffic.

features are *volume-based*, the flow, packet and byte count, and *distribution-based*, the Shannon entropy of source and destination IP and port distributions. These features are commonly used in flow-based anomaly detection.

Time series entries for a five minute period are created by counting the number of flows, summing the number of packets and bytes, and creating map structures of IP address and port distributions. The latter are used to calculate an entropy measure. When all the files that describe services over a five-minute period have been processed, entropy is calculated and all features are stored as one time series entry. The aforementioned baseline time series are created for training data; these are free from anomalies and migration. Conversely, evaluation time series contain anomalies and may contain migration, depending on the experiments we wish to conduct.

As discussed in Section III, wide-area migration can be observed as the removal or addition of new services and associated inbound and outbound traffic at the source and target data centre, respectively. We simulate this behaviour as follows: at the source data centre a set of services is selected for the baseline, and a further set are marked for migration at a given time t. For the evaluation time series the creation starts without any changes. However, when t is reached in the dataset, the time series are created without services that were marked for migration. This simulates the stopping of the migrated services at the source data centre. Meanwhile, for the target data centre, the process is the reverse. When t is reached, the set of services on which the time series are based will further include the set of migrated services.

It can be seen that using this toolchain we can flexibly create evaluation scenarios that include a range of attack behaviours, using the FLAME tool, and service migration activity of different magnitudes. These scenarios can then be provided as input to different flow-based anomaly detection techniques, as discussed in the following section.

V. The Impact of Migration on Detection Performance

In this section, we discuss the affect wide-area virtual service migration has on two anomaly detection approaches: spectral analysis, based on Principal Component Analysis, and clustering using Expectation-Maximisation (EM).

A. *Principal Component Analysis-based spectral analysis*

To carry out the analysis, a traffic data profile is created that is based on the features discussed in Section IV, using Principal Component Analysis (PCA). Anomalies are detected based on the difference of the baseline and evaluation profiles. The anomaly detection results are represented by a Receiver Operating Characteristic (ROC) curve, as shown in Fig. 3.

We conducted a number of experiments by varying the type and intensity of anomalies, and examining the affect of migration at a source and target data centre. Initially, a smaller dataset was used to determine that the expected behaviour of the anomaly detection approach occurred, with different anomaly types and intensities. This proved to be the case. Subsequently, we ran experiments with the top 200 TCP services, whilst using different virtual service migration intensities, calculated as a percentage of the overall traffic. In these experiments different anomaly types were injected – a DDoS attack, and a vertical and horizontal port scan. Fig. 3 show the results from experiments with a DDoS attack on a source data centre, with an anomaly intensity of 150,000 anomalous flows per five minute period, injected into 600,000 to 800,000 flows – this difference in the number of flows relates to diurnal variations in network usage, for example. Hence, the anomalies will account for 20% to 15% of the traffic at the moment of injection. Results are shown for anomaly detection performance without (Fig. 3(a)) and with (Fig. 3(b)) migration. Each plot includes results based on a different number of PCA components, ranging from 1 to 8, and denoted by k.

As mentioned earlier, Fig. 3(a) shows the results from anomaly detection on an evaluation set of data without migration. Anomalies are detected with 100% TPR for $k = 3$ and higher. For the lower dimensionalities of the analysis model, for a 2% FPR, the TPR is 20% and 40% for $k = 1$ and $k = 2$, respectively. This gives us the baseline performance of the PCA-based anomaly detection technique, with which to compare the impact of migration.

We added the migration component, starting by migrating 2.5% of the traffic. The resulting ROC curve (not shown here for space reasons) looks similar to the one for the baseline traffic (Fig. 3(a)); the TPR values for all k are the same.

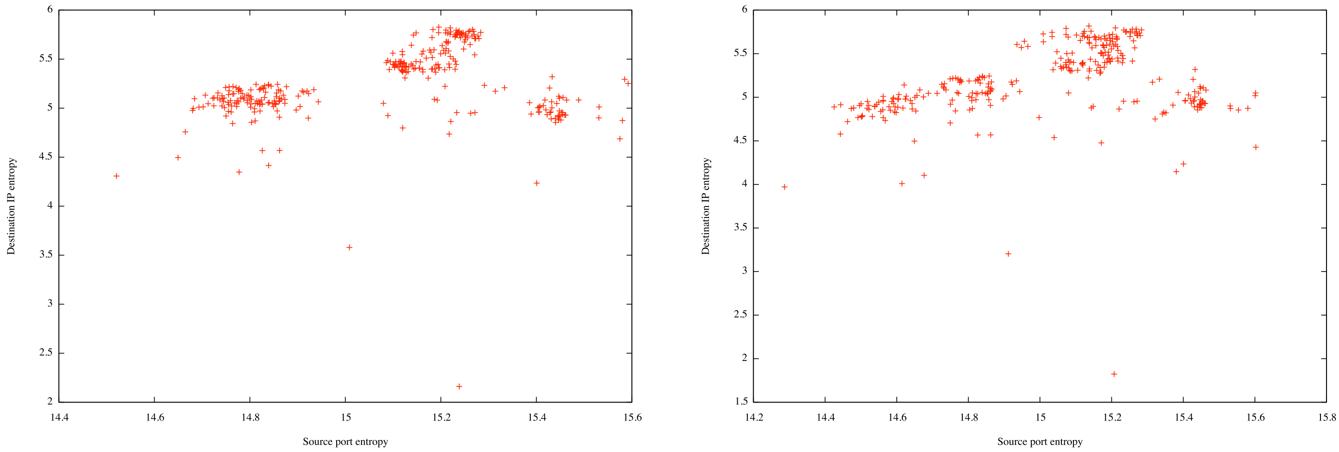

(a) Clustering *before* migration

(b) Clustering *after* migration

Fig. 4. Results of the Expectation-Maximisation (EM) clustering algorithm, plotting source port against destination IP address entropy, before virtual service migration (4(a)), and after (4(b)). A visual shift in clusters can be observed.

For lower dimensions, i.e., for $k < 4$, we observed the anomaly score thresholds differ than for the baseline behaviour. However, we conclude for service migration of this magnitude there is little impact on detection performance.

We conducted further experiments with different magnitudes of migration. For 10% traffic migration, changes in the ROC curves become apparent. For clarity of presentation, we show results from experiments with 20% traffic migration in Fig. 3(b). Migrating this number of services (and its corresponding traffic) could be caused by a single failure in a data centre, or represent an aggregate from several smaller migrations, e.g., due to load balancing. As shown in Fig. 3(b), for $k = 2$, improved TPRs can be observed for selected FPRs, indicated by the shift in the curve to the upper-left, suggesting an improvement in performance, when compared to the results without migration. However, the TPR is still poor in comparison to higher values of k. We suggest the detection algorithm should not be used in this configuration. Perhaps the most interesting observation in Fig. 3 relates to the result for $k = 4$, whereby a noticeable drop in TPR can be observed for comparable FPRs, caused my migration, e.g., for a FPR of 2%, a 20% drop in TPR can be seen because of migration; for lower FPRs this phenomenon is more pronounced.

The results show in Fig. 3(b) show good detection performance after migration for values of $k > 4$. In other experiments that use different experimental parameters, e.g., anomaly type and intensity, not shown here for space reasons, we observed instances when this was not the case. Furthermore, we noted this in experiments we conducted using Tsallis entropy, as opposed to the results shown here that are based on Shannon entropy. As expected, we observed similar behaviour at the target data centre. For a further examination of these results, we refer the reader to [16].

To understand the impact of these results on the usability of the algorithm under migration conditions, one must consider the values that underlie the rates in terms of the number of flows that were incorrectly detected. Based on observations from our dataset, the number of flows that are observed for a week could reach approximately one billion, amongst which there are three million injected anomalous flows. For the experiments shown in Fig. 3, with $k = 4$ and FPR of 2%, the successful attack detection rate drops from 100% to 80%. This means that a fifth of the anomalies are missed as being attacks, namely 600,000 anomalous flows. This potentially represents a significant vulnerability to a large cloud data centre and the services it hosts. If we decide to increase the detection threshold, such that the TPR returns to 100%, the FPR increases to 2.5%. This half percent results in false alarms associated with five million non-anomalous flows aggregated into five minute bins – an unusable number for an operator.

B. *Expectation-Maximisation-based Clustering*

Another approach to detecting anomalous behaviour is based on clustering, which functions by assigning data points that have similar features to cluster structures. Items that do not belong to a cluster, or are part of a cluster that is labelled as containing attacks, are considered anomalous. There are a number of clustering algorithms available; we chose to base our experiments on the Expectation-Maximisation (EM) algorithm, because it negates the need to pre-define a number of clusters, was shown to give good detection performance in previous research, and its availability in the widely-used WEKA machine learning libraries[3].

In these experiments, we extract the time series entries that have an anomaly score greater than zero, after the spectral analysis discussed in Section V-A. This is done to significantly reduce the number of data points to consider, and to focus on the anomalies caused by attack and migration behaviour. Subsequently, the entries are clustered based on the four entropy-based traffic features, source and destination IP address and port distributions, in order to ignore the fluctuations in number of traffic flows, packets or bytes.

Fig. 4(a) depicts the clusters that are formed without migration; source port and destination IP address entropy features are shown. It appears that three clusters can be distinguished, including some more dispersed entries. However, the clustering

[3] http://www.cs.waikato.ac.nz/ml/weka/

algorithm determines nine clusters, with a log likelihood 1.78. Fig. 4(b) shows what happens after migration. The clusters have dispersed because some of the distributions have been changed after migration. Thus, on this plot half of the data is before migration and half after. It appears as if parts of the clusters have shifted to the right. The clustering algorithm distinguishes only seven clusters, compared to the nine before migration. These clusters are larger and more dispersed, which is confirmed by the log likelihood falling to 0.72. This means that the probability that a selected instance will be put in the correct cluster has decreased more than twice. Even if we apply the algorithm with well-formed clusters after migration, the probability to correctly classify an instance remains low. These results, in a similar manner to those obtained using Principal Component Analysis, indicate that virtual service migration makes the use of this form of clustering approach unreliable for anomaly detection.

VI. Conclusions and Discussion

Operators of critical infrastructures are considering moving their high assurance ICT services to the cloud. This implies heightened security requirements; attacks could lead to outages in services that our society depends on, or result in sensitive data being disclosed. An important security measure is to be able to detect when a cloud data centre, and the services that it hosts, are attacked via the network. There are numerous ways to achieve this, including detecting anomalies in network flow summary data – an approach that has seen significant research interest and deployment. One of the essential characteristics of cloud computing is the use of virtualisation technology, which enables services to migrate between different underlying physical infrastructures, both within and across different cloud data centres. Virtual service migration can be used to realise load balancing strategies and improve fault-tolerance to underlying hardware failures, for example. In large cloud data centres, virtual service migration can happen relatively frequently as a consequence of these day-to-day operations.

We have examined the affect that wide-area virtual service migration – i.e., migration between cloud data centres – has on contemporary techniques for detecting anomalies in network flow summary data. We have shown that spectral analysis-based detection, using Principle Component Analysis (PCA), and the Expectation-Maximisation (EM) clustering algorithm can be adversely affected by virtual service migration. We argue that, under certain attack and migration conditions, the number of attacks that are missed and false alarms generated by these techniques could render them unreliable and unusable, respectively.

In search of a solution to this problem, we carried out experiments in which virtual service migration behaviour was incorporated into the baseline "normal" behaviour time series. After all, migration of services is arguably representative of the day-to-day operation of a cloud data centre. We found the results of these experiments did not lead to improved detection performance. Further work will examine the reasons why this did not help. Our current thinking about a solution to this problem involves keeping records of the services that are migrated, and using these in a post-detection processing phase to suppress alerts that relate to service migration. We appreciate this approach is not ideal as it requires maintaining and migrating additional state about a virtual service and its clients, which increases overheads and introduces potential privacy issues. Conversely, another approach might involve correlating the alerts from different data centres, in order to determine whether similar behaviour that is indicative of an attack or other problems have been observed.

Acknowledgements

The research presented in this paper was partially funded by the EC under the FP7 project SECCRIT (Grant No. 312758). The authors are grateful to Dr Bernhard Tellenbach for his support obtaining the PCA-based analysis results.

References

[1] M. Dekker, "Critical Cloud Computing A CIIP perspective on cloud computing services," white paper, December 2012, European Network and Information Security Agency (ENISA).

[2] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 15:1–15:58, Jul. 2009.

[3] B. Grobauer, T. Walloschek, and E. Stocker, "Understanding cloud computing vulnerabilities," *IEEE Security Privacy*, vol. 9, no. 2, pp. 50–57, 2011.

[4] B. Tellenbach *et al.*, "Accurate network anomaly classification with generalized entropy metrics," *Comput. Netw.*, vol. 55, no. 15, pp. 3485–3502, Oct. 2011.

[5] C. Pascoal *et al.*, "Robust feature selection and robust PCA for Internet traffic anomaly detection," in *IEEE INFOCOM 2012*, Orlando, FL, USA, March 2012, pp. 1755–1763.

[6] H. Ringberg *et al.*, "Sensitivity of PCA for traffic anomaly detection," *SIGMETRICS Perform. Eval. Rev.*, vol. 35, no. 1, pp. 109–120, Jun. 2007.

[7] W. Lu and H. Tong, "Detecting Network Anomalies Using CUSUM and EM Clustering," in *Advances in Computation and Intelligence*, Macao, China, June 2009, vol. 5821, pp. 297–308.

[8] I. Syarif, A. Prugel-Bennett, and G. Wills, "Unsupervised Clustering Approach for Network Anomaly Detection," in *Fourth International Conference on Networked Digital Technologies (NDT 2012)*, Dubai, AE, April 2012, pp. 24–26.

[9] D. Brauckhoff *et al.*, "Impact of packet sampling on anomaly detection metrics," in *6th ACM SIGCOMM conference on Internet measurement*, Rio de Janeiro, Brazil, October 2006, pp. 159–164.

[10] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology (NIST), Tech. Rep. Special Publication 800-145, September 2011.

[11] K. V. Vishwanath and N. Nagappan, "Characterizing cloud computing hardware reliability," in *Proceedings of the 1st ACM symposium on Cloud computing*, Indianapolis, Indiana, USA, 2010, pp. 193–204.

[12] M. Nelson, B.-H. Lim, and G. Hutchins, "Fast transparent migration for virtual machines," in *USENIX Annual Technical Conference*. Anaheim, CA, USA: USENIX Association, 2005, pp. 25–25.

[13] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, Melbourne, Australia, November 2010, pp. 267–280.

[14] D. Schatzmann, "Correlating flow-based network measurements for service monitoring and network troubleshooting," PhD Thesis, ETH Zürich, 2012.

[15] D. Brauckhoff, A. Wagner, and M. May, "FLAME: a flow-level anomaly modeling engine," in *Conference on Cyber security experimentation and test (CSET'08)*, San Jose, CA, USA, 2008, pp. 11–16.

[16] K. Adamova, "Anomaly detection with virtual service migration in cloud infrastructures," Master's thesis, D-ITET, ETH Zurich, 2013. [Online]. Available: ftp://ftp.tik.ee.ethz.ch/pub/students/2012-HS/MA-2012-17.pdf