Jonna-Janita Eskelinen

Conducting Risk Assessment
Cloud Provider Perspective

Helsinki Metropolia University of Applied Sciences

Master's Degree

Information Technology

Master's Thesis

11 October 2016

Metropolia

Author	Jonna-Janita Eskelinen
Title	Conducting Risk Assessment: Cloud Provider Perspective
Number of Pages	51 pages + 1 appendices
Date	11 October 2016
Degree	Master of Engineering
Degree Programme	Information Technology
Instructor(s)	Ville Jääskeläinen, Head of Degree Program

The goal of this Master's Thesis was to study the risks a cloud service provider should be prepared to meet and perform a risk assessment for the case company's proof of concept cloud where they act as the service provider. Due to security reasons, the study concentrates on the risk assessment process instead of the specific results, but the results are discussed on high level in order to evaluate the suitability of the selected methods for the cloud.

This thesis was done in two parts. First the previous research on cloud security risks was studied and then the actual risk assessment done. Most of the previous research was made from the viewpoint of the cloud user instead of a provider, but in this study the reports are analysed to determine which of the risks apply also to the cloud provider. The risk assessment performed in this study was qualitative and the framework from the ISO/IEC 27005:2011 standard. STRIDE was selected as the threat modelling method. As a secondary plan for identifying the threats and vulnerabilities, a questionnaire with industry best practices was prepared.

Two workshops were held with the case company during the risk assessment process. The first one concentrated on identifying the risks and the second on rating the risks. In the end STRIDE threat modelling was found to be challenging in a cloud environment and the secondary plan of identifying deviations from the industry best practises provided better results. Performing the risk assessment during the proof of concept phase limited the tools and data available, but the results were found to be valuable as a preparation for the production environment.

Based on the risk assessment findings, the effect of the selected methods on the results were evaluated, a comparison of different threat modelling methods is presented here together with recommendations for future risk assessments.

Keywords	Cloud, Cloud computing, IaaS, Information security, Risk assessment, Security risk, Threat modelling

Metropolia

Contents

List of Tables

List of Figures

List of Abbreviations

3PP	Third Party Product. Products that are developed by another company.
API	Application Programming Interface. Set of tools and protocols used to build applications. Defines how components interact.
CRM	Customer Relationship Management. Technologies and strategies that are used to manage company's customer relations over their lifecycle.
DNS	Domain Name Service. Service used to translate alphabetic names into IP addresses.
DoS	Denial of Service. Attack where the target's resources or networking capacity are depleted by overwhelming the system with traffic or requests.
HTTP	Hypertext Transfer Protocol. Client-server protocol used for data transfer by the World Wide Web.
IaaS	Infrastructure as a Service. Cloud service model where cloud provider offers hardware, storage and infrastructure to the consumer. Everything starting from the operating system is controlled by the consumer.
OS	Operating System. The software that acts as interface between the hardware and software.
PaaS	Platform as a Service. Cloud service model where the cloud provider offers a platform usually for software development.
PoC	Proof of Concept. A prototype implementation where certain method or plan is proven to be feasible. The implementation may be complete or lack some of the features.
SaaS	Software as a Service. Cloud service model where the cloud provider offers an application over the internet.

SLA	Service Level Agreement. Contract between the service provider and customer that defines what services are being delivered and the performance expectations related to them.
VM	Virtual Machine. Emulated machine that runs on shared hardware.
XML	Extensible Markup Language. Markup Language used to store and transport data.

1 Introduction

The use of cloud environments has grown quite common and companies are moving more and more services to the cloud. The benefits of cloud - scalability and cost effectiveness - are hard to ignore. However, cloud environments also bring many new challenges, many of them related to information security. One the biggest concern for companies is the location of the data, which by legislation may have to be in the same country or at least in the same continent. This and also other factors related to security and privacy have lead companies to set up own their clouds.

Any service exposed to the public is always vulnerable to attacks. Some outlines can be viewed in Verizon's yearly Data Breach Investigations Report. The data for 2015 covers 64199 security incidents of which 2260 lead to an actual data breach. Breach in this case means a confirmed exposure of data to an unauthorized party. The rest were incidents where the availability, confidentiality or integrity of data was potentially compromised. [1]

The motives were mostly financial as seen in Figure 1, though it should be noted that the motives often overlap [1].

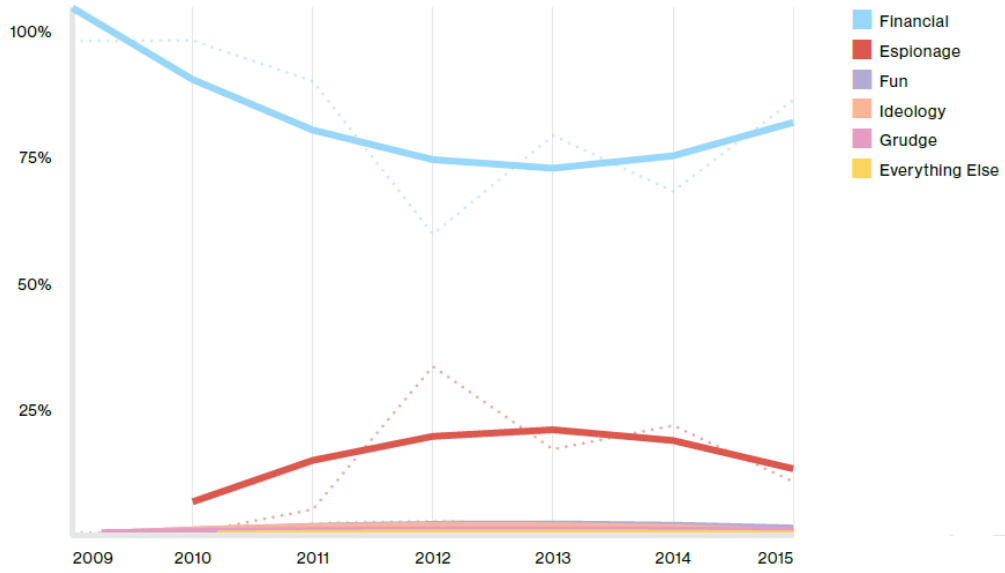

Figure 1. Motives behind breaches [1]

The report excluded thousands of cases with secondary motives where the attacker's goal was to use the system to launch or support consequential attacks like denial of service (DoS) or hosting malware. If included, they would have outweighed the other motives [1]. With cloud based services, the trend can be expected to be the same. A public cloud is probably even more tempting target than a traditional server or s device, since it can be used to launch attacks towards other targets or at least to harness more resources.

When assessing security risks it is not only the attacks that should be evaluated. The risks can be a result of misuse or human errors. The motive does not have to be malicious either. Instead the risks can rise from flaws in operational procedures and usage, which brings a whole new spectrum of threats to consider.

The goal of this study was to investigate the process of making a risk assessment and also evaluate the suitability of the selected methods for a cloud environment. However, the specific results of the risk assessment are out of the scope of this thesis due to confidentiality reasons.

The risk assessment was done for the case company on their proof of concept (PoC) cloud environment where they act as a cloud service provider. The PoC is a prototype system used to provide evidence that the implementation is feasible and the plan works. In this case the PoC cloud did not have all features of a production system, but it was a working cloud environment. The reason for performing the risk assessment at this phase, was to better prepare for the establishment of the production environment. When done already in PoC, the risks discovered could be taken into consideration when designing the final environment. Changes to the system would be much easier to implement before going to production.

Groundwork for this study was done by examining the security risk listings provided by different cloud organizations and communities. The objective was to understand what kind of threats and risks are commonly seen in cloud environments. Most of the risk reports available were compiled from the viewpoint of the cloud users, but in this study they were examined in order to understand what the cloud provider should be prepared for.

The actual risk assessment was conducted during two workshops with the cloud provider. In the first workshop threats and vulnerabilities related to the system components were identified and deviations from recommended security controls were also looked at. In the second workshop the risks' impacts and probabilities were evaluated from the viewpoint of the cloud provider.

The report from the project is divided into six chapters. In Chapter 2 the basic concepts related to cloud and security risks are discussed. The chapter starts by presenting the terminology related to cloud and information security risk and continues introduce the methods selected for this study. Next, the different organizations related to the cloud along with previous research on the cloud security risks are examined. Finally, the chapter provides conclusions on the findings.

Chapter 3 briefly goes through the standards and recommendations related to information technology security, whereas Chapter 4 is dedicated to the risk assessment process. The risk assessment framework selected for this study is based on ISO/IEC 270005 standard, but the chapter goes beyond the framework to briefly discuss the exact methods used in this project.

In Chapter 5 the practical implementation of the risk assessment is described. It includes preparation and the two workshops held with the customer. In Chapter 6 the results of the assessment are discussed. Furthermore, the impacts of the selected methods on the results are evaluated. Finally, in Chapter 7 conclusions on the different methods are presented along with recommendations for future risk assessments.

2 Security Risks in Cloud

This chapter goes through the concepts related to security risks in the cloud. First the terminology related to cloud is clarified and then the basics of information security risk are discussed. Finally previous research on cloud security risks is introduced.

2.1 Definition of Cloud

One of the most referred description of a cloud is by National Institute of Standards and Technology (NIST):

> Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. [2]

Before defining the different aspects of cloud more accurately, the terminology should be clarified. Firstly there is the cloud provider, which is the organization or entity maintaining and owning the cloud hardware and resources. The second is the cloud consumer, or user organization, who buys the resources or service from the cloud provider. The third term often used with a cloud is a tenant, which is essentially the same as a consumer but should not be mixed with one person as a user. Basically one environment inside the cloud system is referred to as a tenant.

NIST lists five central characteristics that separate the cloud from other deployments [2].

- **On-demand self-service:** No human interaction with the cloud provider is required when the consumer needs to provision resources.
- **Broad network access:** Offers availability over the network, but also usability for different kind of clients like tablets or laptops.
- **Resource pooling with multi-tenancy:** Resources like storage or memory are pooled and assigned to different consumers as needed.
- **Rapid elasticity:** The system can be scaled up or down, even automatically, based on demand

- **Measured service:** The resources used can be measured in order to provide data to both consumer and provider. This data is often used as basis for Service Level Agreements (SLA) and billing.

NIST's definition is not the only one for cloud and the wording may slightly vary depending on a vendor or an organization, but usually they all share same characteristics where the cloud is a pool of resources which can be scaled dynamically without interaction with the service provider. A cloud uses virtualization techniques, but virtualization alone does not equal cloud.

2.1.1 Cloud Deployment Models

There are three deployment models available for a cloud: public, private and community. Some also list fourth, a hybrid, which is a combination of any of the three. The deployment model depends mainly on who the cloud consumers are.

A public cloud, for example Amazon AWS or Microsoft Azure, is available for everyone. The computing resources are usually the largest in this deployment model and elasticity the highest. The communication towards the cloud is done via the public Internet and the hardware is located on the cloud provider's premises [2].

From the security point of view, the public cloud is the most attractive model to the attackers, since the access is not limited for a certain group. A large consumer base means there's more likely to be a malicious user included. Consequently, multi-tenancy and shared infrastructure pose additional risks with a public cloud. The resources used by an organization are separated only by policies implemented in the cloud provider's software. Vulnerabilities or mistakes in the access controls or work procedures could potentially expose organization's data to unauthorized party. [2][3]

Public clouds can also introduce a problem with the location of the data. The consumer is not able to verify where the data is stored or how it is destroyed when no longer needed. This can cause legal issues in some cases. However, location can also be a benefit. A public cloud can offer services at a lower cost, since the datacenters can be built in places where either the costs for maintaining or establishing the infrastructure are cheaper. [2][3]

The second model, a private cloud, is used by a single company or organization. The owner and manager of the infrastructure can be the organization or an external party. In both cases the company must invest in the infrastructure, and have or buy the competence to maintain the cloud system. The hardware itself can be located either on or off premises, but it is not shared with other organizations. [2][4]

Unlike with public cloud, the location of the data is not usually a concern with a private cloud. The organization knows and is able to verify where the data is physically stored. However, additional investments are required if the data has to be stored in geographically separated locations. Compared to public cloud, implementing geo-redundancy is much more expensive with a private cloud. In general, downside of a private cloud is that it will not have as large resources disposable as a public cloud and will not reach the same elasticity. [3]

When thinking of security, a private cloud has the same security issues brought on by shared technologies. A malicious insider can do as much harm in this model as with a public cloud, though the attacker base is smaller. With private cloud the data might not end in the hands of a competitor, but for example financial data can be highly sensitive even if leaked inside organization. [2][3]

The third option, a community cloud, is close to a private cloud, except it is shared by multiple organizations with similar interests and common goals. Usually the policies and requirements of the organizations involved in the community cloud are also compatible [3]. The benefits and downsides or the model are usually closer to a private than a public cloud, but the other consumers in the cloud are considered trusted [4].

From the security point of view, community cloud has a smaller attacker base than a public cloud, but the same risks with multi-tenancy exist apply to all deployment models. Whether the community cloud is located on one of the organizations' premises or outsourced, the other consumers must access the system either via public Internet or a leased line. This rises some extra security concerns compared to private cloud which could be on-premises for the organization. [3]

The last model is a hybrid cloud which can be any combination of a public, private and community cloud. The clouds are their own entities, but they are connected by technol-

ogy that allows data portability. The benefits and security issues will depend on the combination chosen, but the resulting cloud system can be complex to map or characterize. [3]

2.1.2 Cloud Service Models

Besides the different deployments models, a cloud also has three service models that define what kind of services the cloud consumer uses from the provider. The models are Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS).

The first, Infrastructure as a Service, offers only infrastructure. Two examples of IaaS providers are Amazon AWS and Microsoft Azure though both offer also other service models. With IaaS the cloud consumer does not purchase physical servers, networking or storage devices, but instead buys them as a virtualized resources from the cloud provider [4]. With IaaS user can choose from a selection of storage options and has control over the operating system (OS), but they have no access to the hardware. Figure 2 displays how the control is divided between the provider and the consumer.

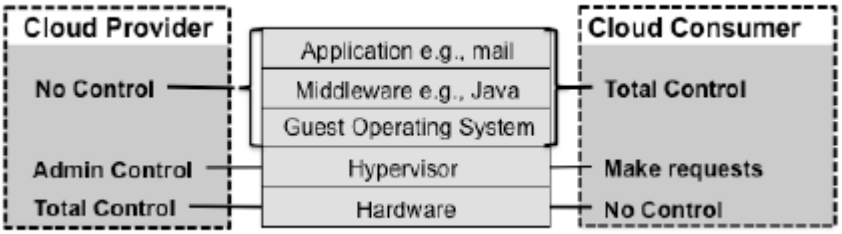

Figure 2. Scope of control in IaaS by NIST [3]

As seen in Figure 2, the consumer is able to make requests towards the hardware via a hypervisor layer. Simplified, the hypervisor is software that uses hardware to produce virtual machines (VMs). From the consumer point of view they can access the VM very similarly to having a physical machine. They can configure and power the machine on or off and even modify the network, but in the background it is the hypervisor that handles all the actions. The hypervisor layer takes care and polices that required resources are available before executing any actions. [3]

Compared to other service models, the consumers have the most control over the environment with IaaS. Consequently they also have to take greater responsibility of the security controls implemented. Everything above the hypervisor layer is handled by the consumer, meaning also updates and security configurations for the OS have to be administrated by the consumer. [3]

The layers might not always be as strict as displayed in Figure 2, but it gives a good overview. For example, the cloud provider may only offer a selection of operating systems and the consumer can have some interfaces to the storage or networking but not full control of either.

The second model is Platform as a Service (PaaS), where the user does not have control over the OS or the underlying infrastructure. The consumer is not able to create virtual machines. Instead only a platform with specific, predefined tools is offered, usually for software development [3]. An example of PaaS provider is Google App Engine, but also Microsoft Azure and Amazon AWS can be used as PaaS. Figure 3 shows how the control is divided between provider and consumer in PaaS.

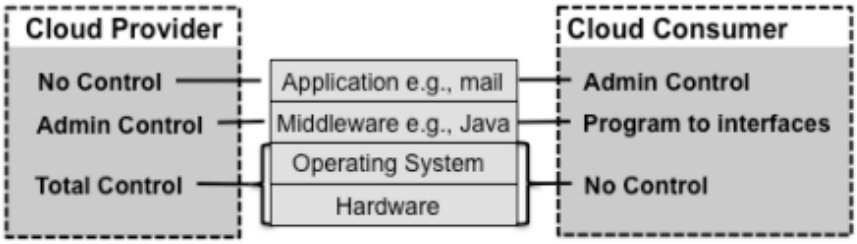

Figure 3. Scope of control in PaaS by NIST [3]

As seen in Figure 3, the provider only offers programming interfaces towards the middleware for the consumer, who then uses them to build an application. Basically the consumer can develop and host a web application without having to invest in the physical infrastructure. The consumer may have the possibility of affecting configuration settings on the platform in some cases, but all layers below are handled by the provider [2] [3].

One of the benefits for the consumer is cost-savings, since there is no need to build a whole infrastructure for the development. It can also be practical for students or people learning software development. In a larger scale, the problems related to platform are

also migrated to the provider and the consumer does not have to worry about the scalability. The billing models may vary depending on the consumer, but in large cloud environments the costs are fairly low. Usually the fees are based on resources consumed or a used platform time [3].

The last option is Software as a Service (SaaS). With SaaS the application, for example an email, is hosted in the cloud and a web interface to it is offered. The consumer is not able to make changes to the application, except maybe some small configuration settings. In this model, the user does not have access to any operating system or even a platform [2] [3]. One of the most known SaaS solutions is Salesforce Customer Relationship Management (CRM).

The benefits of SaaS are largely the same as for PaaS. The consumer does not have to worry about scalability, infrastructure or management of the application. As a downside, the security in this case is also out of the hands of the consumer. The consumer must be able to trust that the security of the provider is sufficient for the use. [3]

In this study the case cloud being evaluated is offering Infrastructure as a Service. In the future, the company will also offer PaaS and SaaS, but at the time of the thesis, these features were still being developed.

2.2 Information Security Risk

The risks and benefits of cloud models were already briefly mentioned in the previous sections, but before going into the details, the concept of information security risk is discussed. The importance of understanding the information security risks rises from the fact that it is not possible to make a system or service completely secure while maintaining usability. The purpose of identifying the risks does not necessarily mean preventing the security incidents from happening. In some cases the costs of preventing an incident may be higher than the value lost in the actual event. Investigating the risks makes it possible to choose the most cost-efficient option between accepting, mitigating or trying to avoid a risk.

Typically an information security risk refers to the possibility that the confidentiality, integrity or availability of the information system is compromised. In the simplest form risk is defined as:

Risk = Likelihood * Impact

The likelihood in the formula indicates how probable the security incident is and the impact represents the effects of the incident to the organization. [5] [6]

Information security risk cannot be discussed without the concepts of threat and vulnerability. ISO/IEC 27000:2012 standard defines threat as “a potential cause of an unwanted incident, which may result in harm to a system or organisation”. Vulnerability on the other hand is defined as “a weakness of an asset or control that can be exploited by one or more threats” [7]. The relation between threat, vulnerability and risk is displayed in Figure 4.

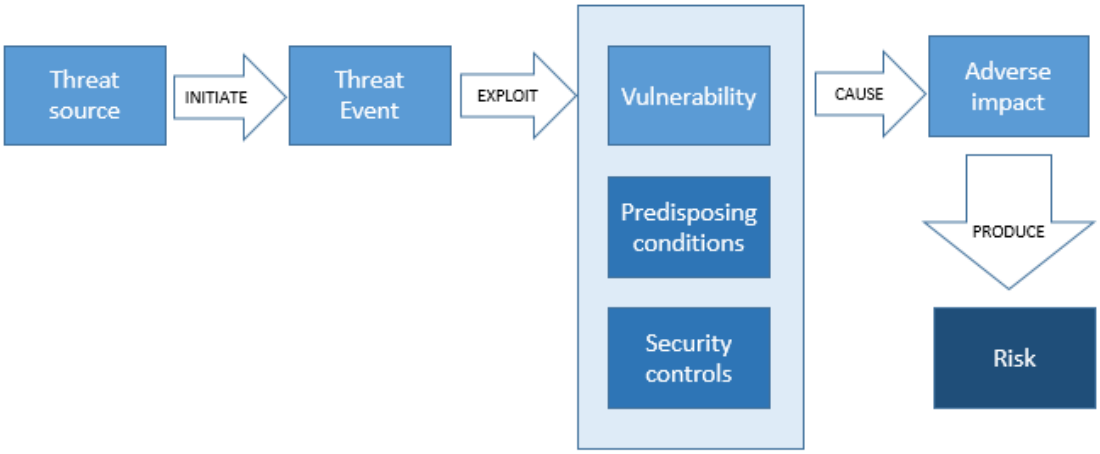

Figure 4. Risk model with key risk factors [5]

Risks emerge when there is a possibility that a threat will exploit a vulnerability and cause harm to the organization [7]. The likelihood in the risk formula presents itself twice in the risk model. Firstly, there must be a likelihood that the threat source initiates the threat event either accidentally or on purpose. It requires that there are means and in some cases motive. Secondly, there is the likelihood of success for the exploit. This is affected by, not only the vulnerability, but the predisposing conditions and security controls of the system. For example, a geographical location is a predisposing condition if considering a risk for an earth quake, whereas weak structures can be a vulnerability. If the threat has a probability for successfully exploiting the vulnerability, the resulting impact together with the likelihood produce the risk. [5]

The evaluation of impact and probability and their relation to a risk assessment process will be discussed more in detail in Chapter 4, but as seen from the risk model, there is no risk without a threat or a vulnerability. Therefore, one of the key factors in handling information security risks is understanding the threats. The following two sections briefly introduce strategies for finding the threats.

2.2.1 STRIDE Threat Modelling

STRIDE is a threat model by Microsoft, published in 1999 by Loren Kohnfelder and Praerit Garg. The acronym is derived from the following words: spoofing, tampering, repudiation, information disclosure, denial of service and elevation of privilege. The terms are the opposites of security properties which all systems attempt to implement: authentication, integrity, non-repudiation, confidentiality, availability and authorization. [8].

The STRIDE model reminds to think about those six threats when investigating what can go wrong in the system. It is not important which category each threat belongs to, instead the purpose is to offer ideas for locating the threats. Hence, the different angles can be applied to multiple situations. They can be used when examining network, system or even operative functions.

The first term, spoofing, refers to pretending to be someone or something and it is the opposite of authentication. In the simplest form it can be calling a person and pretending to be from help desk. Spoofing is one of the most common threats that appear in social contexts, but it applies to technical systems equally. For example, it can be a program that is named to look like something else or a forged DNS entry that takes the user to a malicious web page when trying to access a bank's site. [8]

The second term, tampering, means modifying information or a system without permission. Examples are manipulation of HTTP headers or network traffic. However, it could also mean editing a customer database or a system configuration file. In those cases integrity of the system is being compromised and the information can no longer be guaranteed to be accurate or complete. [8][9]

The third threat, repudiation, takes place when something is done and there is no way to prove it happened. It is the opposite of security property called non-repudiation. For example, signing a document is a proof of non-repudiation, meaning that one cannot deny

they have read it. In information technology repudiation is usually associated with logging. Consequently, saving logs outside a system makes it harder for the attacker to hide their tracks by modifying or erasing the logs. [8][9]

The fourth in the list, information disclosure, is a breach of confidentiality where information is being seen by a party that should not have any access to it. Eavesdropping a network communication is one example of information disclosure. Similar effect can be the result of having incorrect permissions on a file [8]. In the cloud world information disclosure is one of the biggest concerns for the cloud consumers. When a user uploads their data to the cloud, there has to be sufficient safeguards to preserve the confidentiality of the information during transfer and also in storage.

Another important factor with cloud environments is availability. If the consumer is running applications or other business critical systems in the cloud, it is vital that the service is always accessible. Denial of service threats jeopardize that. The threat is often associated with network attacks, where the bandwidth of a server or a link is exhausted, but it can also mean overprovisioning system resources like CPU or memory until the system becomes unusable [8]. Another cloud-related example is overflowing an application programming interface (API) hosted in a cloud. If the cloud resource is being changed based on API calls, such attack could exhaust the credit limit assigned for the system.

The last term, elevation of privilege, happens when normal user is able to perform administrative actions that they should not be able to execute. The entity does not have to be a user, instead it could also be a process. In the programming context, this can happen if the authorization checks are not sufficient. Consequently, the threat is considered to be the opposite of authorization. [8][9]

However, the borders of each category are not always clear. Repudiation can happen via tampering of data or information disclosure by spoofing. When modelling threats it is not important that the right category is chosen. The main point is that the threats are found and listed. Similarly, probability of the threat is not a concern at this stage. The likelihood of the event will be evaluated in a different stage of risk assessment.

2.2.2 Attack Libraries

Another method of discovering threats are attack libraries. Attack libraries are prebuilt collections that include descriptions about cyber-attack patterns. Patterns in this case means detailed technical explanations about the attacks. As opposed to STRIDE, attack libraries are not high level or abstract, instead they have practical examples of attacks that have been done it the past. [8] [10]

One of the largest attack libraries is MITRE's Common Attack Pattern Enumeration and Classification (CAPEC). The library was originally established by U.S. Department of Homeland Security in 2007, but nowadays it is being updated by public participation. The most recent version 2.8 released in December 7 2015 included 504 different attack patterns. [10]

CAPEC list is very detailed and it can be browsed online or downloaded in XML format. The information on each attack contains summary, attack prerequisites, typical severity, likelihood, attacker skill required, resources required, mitigations, related weaknesses and many other factors related to the attack. The security controls against the attack are also well described. [8][10]

The CAPEC library can be searched with keywords, for example "SQL", or by IDs, but it is also organized in domains of attack and mechanisms of attack. The domains are social engineering, supply chain, communications, software, physical security and hardware. For mechanisms of attack there are 17 different categories [10]:

- gather information
- deplete resources
- injection
- deceptive interactions
- manipulate timing and state
- abuse of functionality
- probabilistic techniques
- exploitation of authentication
- exploitation of authorization
- manipulate data structures
- manipulate resources

- analyse target
- physical access
- execute code
- alter system components
- manipulate system users.

The challenge in using CAPEC is the extent of the library. It can be time consuming to go through the whole library if not already familiar with most of the content. Other element to consider is the precision of the patterns. For a less experienced security practitioner the exact details offered by CAPEC can be very helpful. For more experienced professional, the library can also be a limiting factor compared to STRIDE which only offers high level guidance. [8]

In addition to the attack library MITRE also offers similar database for vulnerabilities. The Common Vulnerabilities and Exposures (CVE) is a list of publicly known cybersecurity weaknesses. Similarly to CAPEC, it can be searched with IDs or with keywords. CVE includes over 300 products and services and therefore the keywords used can even be product or component names. The database is used by many vulnerability scanners in order to provide a common lexicon. [11]

2.3 Previous Research on Cloud Security Risks

As cloud services have become more common, studies have been conducted specifically on cloud security risks. The risks vary greatly depending on what kind of services are being hosted and therefore it is difficult to give a comprehensive list that would cover all use cases. However, there are certain organizations that have collected information about risks commonly seen in cloud environments and ranked them accordingly. These listings give a good overview of what kind of issues the cloud service users, and consequently also cloud service provider, should be prepared to meet. These five listings are introduced in the following sections.

2.3.1 CSA - The Treacherous 12

Cloud Security Alliance (CSA) is a non-profit organization established in 2009. They are dedicated explicitly for cloud security. The organization has many large corporate members, such as Microsoft, VMware and EMC, as well as individuals. It provides research,

conferences and also certifications related to the field. It is responsible for a CSA Security, Trust & Assurance Registry (STAR), which is a notable certification in cloud security. [12]

In February 2016 CSA released a report sponsored by Hewlett-Packard that lists the top 12 threats in the cloud. From the five different cloud security reports introduced here, this is the most recent. The study was conducted by first selecting 20 issues and then requesting each working group member to rate their importance to their organization. About 270 respondents took part in the survey. [13]

The results from CSA's report include a description of the risk, a list of the cloud service models that it applies to and a reference to security guidance for each risk. The report also specifies which STRIDE threat category the risk is associated with and offers links to articles with information on similar previous attacks. Therefore the information offered by CSA's report can be directly associated with real-world incidents. [13]

Table 1 lists and describes the risks ranked according to severity as seen by the organizations [13].

Table 1. The Treacherous 12 - CSA's Cloud Computing Top Threats in 2016

Risk [13]	**Description [13]**
1. Data breaches	Confidentiality of data is compromised
2. Insufficient Identity, Credential and Access Management	Insufficient or ineffective authentication methods. Risk that credentials or cryptographic keys are exposed
3. Insecure Interfaces and APIs	Cloud UIs or APIs are not protected well enough
4. System Vulnerabilities	Confidentiality, integrity or availability or the system is compromised due to other tenants, system vulnerabilities or insufficient logging
5. Account Hijacking	Cloud account stolen or hijacked. Company reputation and data compromised
6. Malicious Insiders	Insider in the cloud compromises the data, either intentionally or due to negligence
7. Advanced Persistent Threats	Risk of organized, technically advanced attackers targeting the system, for example government entity
8. Data loss	The data stored in cloud is lost, not just compromised

9. Insufficient Due Diligence	All consequences of moving a service or data to the cloud are not understood and the cloud provider processes do not match the needs or the organization
10. Abuse and Nefarious Use of Cloud Services	Risk that the cloud service is used for malicious purposes, e.g.to launch attacks or to exhaust the cloud's resources
11. Denial of Service	Loss of availability due to denial service attack towards the cloud
12. Shared Technology Vulnerabilities	Risk that the isolation of tenants is not sufficient in the cloud. Fear that the shared resources, e.g. hypervisor is breached

Most of the risks listed in Table 1 apply to both cloud provider and consumer. According to the report, the risks can be seen in all service models, hence the top risks are the same for IaaS, PaaS and SaaS environments.

Even though the greatest risk, data breach is listed as a separate risk, most of the other risks are also related to information getting disclosed to an unauthorized party. Only the risks associated with nefarious use of cloud services, denial of service and data loss do not threaten the confidentiality of the data. With complete data loss, confidentiality is no longer an issue since the information is lost for everyone involved. [13]

When examining the risks from the cloud provider's perspective, it is not only the cloud consumer's data that can be compromised. The provider also has their own system and billing data that needs to be protected. Similarly, insecure APIs or UIs can compromise the whole cloud and not just a specific client. Even if the risk would impact a single tenant, the reputation of the cloud provider would suffer in the event of an incident. Therefore there are no risks that cloud be defined as consumer's responsibility alone.

In addition to confidentiality, another important factor for the cloud provider is the availability of the system. The whole concept of cloud service relies on the fact that the provider can offer near constant uptime for the system. Besides traditional denial of service, also system vulnerabilities, data loss and abuse of cloud services put the availability of the system at risk. Less obvious risks that compromise availability are insufficient identity and access management as well as account hijacking [13]. In cases where it is the provider's own credentials that have been exposed, the result can be catastrophic for the whole cloud infrastructure. The damage that can be done to a system with administrative

credentials is extensive, not to mention the time or cost used to changing all compromised credentials.

Regardless of the direct effect on the provider, it is in their interests to offer safeguards against all of the risks listed. However, depending on the service model, it can be up to the consumer whether they implement the security controls or not. Only the risk related to due diligence can be difficult for the provider to cover, since it greatly depends on the service or data that the consumer will transfer to the cloud. Fully understanding this risk is therefore the responsibility of the cloud consumer. Nevertheless, the provider should try to offer security controls that enable the consumer to fulfil the legal obligations they have.

2.3.2 ENISA – Cloud Computing Benefits, Risks and Recommendations for Information Security

European Union Agency for Network and Information Security (ENISA) is EU funded agency working in the field of network and information security. Established in 2004, they make recommendations and host events that concern European Union's laws and policies about information security. [14]

In 2009 ENISA released a risk assessment related to cloud computing. The assessment was done in co-operation with government organizations, academic parties and industry. The assessment includes 35 risks identified and also recommendations related to the cloud and is therefore the most comprehensive one. [15]

Table 2 lists the eight most important risks found in the report. These risks are not in any order of importance. The full list can be found in Appendix 1.

Table 2. Cloud computing risks by ENISA

Risk [15]	**Description [15]**
Loss of governance	Cloud consumer loses control of many aspects related to security
Lock-in	Risk that once data is stored in cloud, it is not feasible to transfer it back. Consumer is locked in with one provider.
Isolation failure	Shared resources and other tenants compromise the user's system

Compliance risks	Cloud provider can't prove compliance to certain requirements and the audit possibilities by the customer are limited
Management interface compromise	The management interface provided by the cloud is not secured sufficiently
Data protection	The security of the data is compromised either during storage or transit
Insecure or incomplete data deletion	The data deletion is not complete since it may not be possible to effectively wipe the disks due to sharing
Malicious insider	Malicious insider in the cloud provider can compromise the data or system

Although the risks listed in Table 2 were identified in 2009, many still correspond to the risks identified by CSA in 2016. The risk with compliance and loss of governance can be covered with CSA's risk related to due diligence, however, incomplete data deletion is not on CSA's list. Likewise, the risk of lock-in with a certain provider is no longer seen as a top risk. This is probably because presently there are more cloud providers to choose from, and therefore it is possible to select one that can guarantee that the customer is able to reclaim their data if needed. The decision to migrate to cloud is typically not taken lightly and consequently the major concern with data lock-in is how to recover the data in case the provider goes bankrupt instead of planning for possible provider change in advance.

However, ENISA's risks listed in Table 2, are the cloud-specific top eight. In addition, the complete listing includes eleven risks that also apply to traditional IT environments and those were not rated against the top eight. These risks are [15]:

- network breaks
- network management (network congestion, misconnection, non-optimal use)
- modifying network traffic
- privilege escalation
- social engineering attacks
- loss or compromise of operational logs
- loss or compromise of security logs
- backups lost or stolen
- unauthorized access to premises
- theft of computer equipment

- natural disasters.

The impacts of these risks can be considered to be higher in cloud than in traditional systems, at least for the cloud provider. From the provider's perspective some of the cloud-specific risks are in fact less important than these. Loss of governance, data lock-in and incomplete data deletion are risks that the cloud consumer must evaluate. The provider can, in some cases, address these risks on certain level in their service level agreements (SLAs), but transferring these risk completely to the provider is usually not possible. Likewise, risks related to compliance depend on the information stored and usage of the cloud. The cloud provider may not get the customer if they cannot fulfil the requirements, but for example payment card industry (PCI) regulations will be extremely difficult to satisfy. Even attempting to meet such a risk is not worth the investment.

2.3.3 NIST - Cloud Computing Synopsis and Recommendations

National Institute of Standards and Technology (NIST), U.S government controlled agency, released a Special Publication 800-146 "Cloud Computing Synopsis and Recommendations" in 2012. The guideline is directed for Federals agencies, but NIST's recommendations are followed by many commercial organizations. The document introduces the benefits and issues associated with cloud environments and also gives recommendations for best practices. The publication discusses IaaS, SaaS and PaaS models in separate chapters, giving distinct observations on each. [3]

Even though NIST SP 800-146 is not a risk listing similar to the ones introduced earlier, the issues raised in the document can be translated to risks. As this study focuses on IaaS, the issues introduced in Table 3 describe the risks found with that service model. [3]

Table 3. Risks related to IaaS cloud

Issue (risk) [3]	**Description [3]**
Network dependence	Risk that availability of the system is not high enough since public Internet is used for the connection
Browser-based risks	Customer's contaminated browser compromises the whole system or the communication is not secure enough

Compatibility with legacy security vulnerabilities	Legacy software run in the cloud by a consumer compromises security
Virtual machine sprawl	Inactive VM's software is not updated accordingly and compromises system security
Verifying authenticity of an IaaS cloud provider web site	Risk that the cloud provider web site's identity is not validated and instead connection to impostor is established
Robustness of VM-level isolation	Risk that the virtual environments are not isolated properly. Malicious tenants could have access
Features for dynamic network configuration for providing isolation	Risk that the networks of each tenant are not isolated completely
Data erase practices	Data deleted is not completely wiped. Backup policies and replication may also complicate deletion

The risks listed in Table 3 are somewhat more definitive than the ones found in the other listings, but they address similar concerns. As the list is directed at the cloud consumers, there are issues that the provider cannot address. From the cloud provider's perspective, most of the risks that apply are related to multi-tenancy and shared technologies. Outdated virtual machines and isolation of machines and networks all fall under that category. In addition, risks to service availability the provider must typically address, since availability is one of the main benefits offered by the cloud.

Unlike the other studies, NIST SP 800-146 also discusses the promises and limitations of cloud provider policies. Typically the providers promise availability that is above 99.5% and offer a refund if the availability is not reached. However, the refund covers only the service fees and not the potential lost business value. Furthermore, it is the customer's responsibility to report about the outage they experienced. [3]

Depending on the provider, policies also include clauses about the preservation of data. Some providers promise to save the data for approximately one month after the consumer has stopped using the cloud, but others recommend backing up the data either locally or to another provider. Most providers reserve the right to resign that obligation if the consumer violates the provider's acceptable use policies. [3]

Usually cloud providers also guarantee that the customer data is not handed out to anyone without legal request. However, many service agreements state that the provider is

not responsible for security breaches or service interruptions caused by malicious activity. Instead the security risks are often placed on the consumer [3]. Nevertheless, that does not mean that the provider should not address the risks related to security. It only means that the provider cannot understand the risks related to each consumer's specific service and therefore the customer should evaluate those themselves.

2.3.4 OWASP – Cloud Top 10 Security Risks

The Open Web Application Security Project (OWASP) is a non-profit, international organization established in 2001. Their mission is to collect and share information on software security in web applications. They are vendor and product independent while promoting the use of standards. [16]

One of the projects of OWASP is "Cloud Top 10 Security Risks", which lists and ranks risks commonly seen in public and hybrid clouds. The list is collected based on input from security professionals and reports from cloud providers. The listing is targeted for companies hosting services in the cloud but also for cloud service providers. [17]

OWASP's ranking is the least formal reference used in this study, as it is yet to be updated beyond pre-alpha release. The first draft was published in 2009 and it has never reached full maturity [17]. However, there are not many publicly available cloud-specific studies and OWASP has a history of producing high quality security guidelines. Most notable listing being "OWASP Top 10" dedicated for web application security risks. The web application listing is quoted by many books and organizations, providing reference for OWASP's competence in the security field. Therefore the "Cloud Top 10 Security Risks" report (see Table 4) was deemed to be reliable enough to be compared with other similar studies.

Table 4. OWASP – Cloud top 10 security risks

Risk [17]	**Description [17]**
1. Accountability and data ownership	Control over the data is lost. Risks related to management of backups, storage location, data deletion
2. User identity federation	Risk of user identity management getting too complicated. Federated identity should be usable over cloud providers

3. Regulatory compliance	If cloud does not adhere to the same regulations or policies as the organization regulatory compliance is lost.
4. Business continuity and resiliency	Responsibility of the business continuity gets transferred to cloud provider. Risk that the business continuity plans do not meet requirements
5. User privacy and secondary usage of data	User's privacy is violated. Secondary usage of data, for example based on behaviour or clicks, could be allowed
6. Service and data integration	Loss of confidentiality during data transfer. The data-in-transit between the customer to the cloud provider should be encrypted
7. Multi-tenancy and physical security	Risk that another tenant in the cloud compromises the confidentiality, integrity of availability of the user's data
8. Incidence analysis and forensic support	Logging by the cloud provider is not sufficient in the events of security incident. The logs of multiple customers may reside in same hardware and therefore be included in investigations
9. Infrastructure security	Cloud provider systems and networks are not configured according to best practices
10. Non production environment exposure	Risk that the user's non-production environment is not configured according to security policies due to not being in official production

The risks listed in Table 4 include many of the same concerns as the listings earlier. For the provider accountability and data ownership may not seem directly applicable, but securing the provider's own configuration and billing data is their responsibility alone. Regulatory compliance as well as secondary usage of data, are not risks that the provider should be prepared to evaluate. Nevertheless, offering sufficient and reliable identity management is the basis for billing and also an important component in tenant isolation. Whether or not the identity federation conforms to the consumer's requirements is a factor when the service contract is signed, but not an actual risk.

The most notable difference with other listings is the inclusion of non-production environment exposure. However, the provider's system should be designed in a way that a single compromised tenant would not endanger the whole cloud. Therefore the threats related to this risk are already included in the infrastructure security if viewed from the provider's perspective.

2.3.5 Conclusions from Risk Listings

The risk listings provided by the four different organizations share many similar issues. This section compares the reports. The risks were grouped according to the most recent study by CSA.

Certain risks were removed from the composite list, because they can only be evaluated by the cloud consumer. Those risks require in-depth knowledge on the consumer's services, or in some cases understanding on a field-specific legislation, and therefore they cannot be included in any general risk assessment by the provider. Such differences between the responsibilities of the provider and consumer were discussed in the previous sections. The results of the comparison are displayed in Table 5.

Table 5. Composite of risks listings

Risk	**CSA [13]**	**OWASP [17]**	**ENISA [15]**	**NIST [3]**
Shared technology vulnerabilities	yes	yes	yes	yes
System vulnerabilities	yes	yes	yes	yes
Data breaches	yes	yes	yes	(yes)
Denial of service	yes	(yes)	yes	(yes)
Data loss	yes	yes	yes	(yes)
Insecure interfaces and APIs	yes		yes	yes
Insufficient identity, credential and access management	yes	yes	yes	
Account hijacking	yes		yes	(yes)
Malicious insiders	yes		yes	
Incidence analysis and forensic support		yes	yes	
Abuse and nefarious use of cloud services	yes		(yes)	
Advanced persistent threats	yes			

It is noteworthy that the four risk studies only included top ranked risks, meaning that the complete lists of risks found in by each analysis could overlap even more. Since such

data is not publicly available, this comparison will give an overview of what is seen as most critical. However, as ENISA provided a list of 35 risks in addition to the top eight, the full list was used in that case.

In the Table 5, it can be seen that the risk related to shared technologies, or tenant isolation, is ranked highest. The risk is distinct for the cloud and consequently it is no surprise all studies included it. The existence of this risk does not depend on the deployment model or even service model of the cloud. Therefore all cloud providers must evaluate and use the necessary security controls to lower the probability of the risk.

Closely related to the same issue are system vulnerabilities. However, as opposed to the shared technologies, this risk is more general. The object of system vulnerabilities is the same libraries and applications that are used in traditional IT environments. As an example, every regular web or email server must address the risks related to these vulnerabilities regardless of being hosted in cloud or not.

All of the studies also mention data breach and complete loss of data in one form or another. In NIST SP 800-146 neither is directly listed as an issue, but the publication includes recommendations for analyzing the data protection mechanisms before committing to any provider.

The last risk that can be said to be included in all studies is denial of service. The risk compromises the availability of the service and in that form it can be seen in all four reports. OWASP's ranking does not directly address the risk, but the risk related to multitenancy covers scenarios where other tenants jeopardize the availability of the system by consuming excessive resources and effectively creating denial of service situation. The risk differentiates from NIST's risk definition where the availability is compromised due to to public Internet, whereas studies by ENISA and CSA include both cases as separate risks. However, since availability is one of the key benefits of cloud, denial of service risks must be taken seriously regardless if they happen in the network or inside the cloud system. Hence the risk is seen as a separate issue in the composite list.

Three of the four reports include risks related to the security of the cloud interfaces and identity management. These risks the cloud provider will also have to address. Account hijacking is also mentioned on certain level in three reports. NIST SP 800-146 lists this risk in the form of consumer having to validate the authenticity of the cloud provider's

web page. If they access an impostor instead, there is a possibility of compromised credentials.

One risk that is not included in the composite list, due to being mainly the concern of the consumer, is incomplete data deletion. Three of the four studies listed this risk in their analysis. However, there is practically no way the cloud provider can guarantee complete erasure of data, since the physical disks are shared by multiple consumers. At most, this can be addressed in the cloud provider's policies.

Lastly, there is a risk that was only included in two of the studies, but which is important for the cloud provider. Malicious use of cloud resources appears on CSA's report, and in a way in ENISA's. ENISA lists the risk from the consumer's perspective as loss of reputation due to co-tenant activities. However, the threat in both cases covers situations where the cloud is used to launch attacks to the outside world. For the provider that has a negative effect on their reputation, but it also consumes cloud resources. Furthermore, the attacks can result in organizations or service providers blocking the cloud provider's IP address range, thus making the cloud services unusable.

3 Information Security Controls

As introduced in the previous section there are a multitude of risks that the cloud provider must be prepared to meet. Section 2.2 introduced the key factors in risk models where the probability of a successful exploit is affected by both the vulnerabilities and security controls in the system. Securing an IT environment is not a simple task, but there are many guidelines that can be followed. Different organizations offer best practices and procedures that help with security. Understanding the security controls also helps to understand the vulnerabilities and risks in the system. The controls may not all be applicable directly to a cloud, where machines and services are regularly deployed and deleted, but most of the practices still hold.

3.1 ISO/IEC 27001:2013

ISO/IEC 27001:2013 is an international standard published by International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) in 2013. The standard lists requirements for Information Security Management System (ISMS), which is used to analyse and address the risks in an organization. It is a formal specification and does not offer specific controls. The standard defines processes for leadership, planning, support, operation, performance evaluation and improvement of the ISMS. It does not detail specific methods which organizations should use to accomplish the requirements. [18]

Since the standard does not offer any specific controls, the details are not introduced here. The relevance with this study comes from the requirements for performing continuous risk assessments and therefore confirms the need for this assessment.

3.2 ISO/IEC 27002:2013

ISO/IEC 27002:2013 is another international standard updated in 2013. The standard has a history of 30 years and it has been updated regularly during that time. It is dedicated to information security in a broad scope, including intellectual property and knowledge, not just computer data. Unlike ISO/IEC 27001, this standard defines a set of specific security controls that should be applied. ISO/IEC 27002 covers the following 14 sections [19]:

- information security policies
- organization of information security
- human resource security
- asset management
- access control
- cryptography
- physical and environmental security
- operations security
- communications security
- system acquisition, development and maintenance
- supplier relationships
- information security incident management
- information security aspects of business continuity management
- compliance.

Under each of the topic there is a list of detailed controls, totaling in 114. The controls are specific, but do not dictate the technologies or exact methods used to implement each control. It instructs for example to perform event logging, protect against malware and segregate networks [19]. All those and many more can be directly applied to the cloud. Neglecting to implement the controls will result in increased risks to the system. Therefore knowing this standard will also help in evaluating the risks.

3.3 CIS Critical Security Controls

Center for Internet Security (CIS) is a U.S non-profit organization established in 2000. The mission of the organization is to identify and provide best practices in cyber security for both the public and private sector. The organization is an internationally recognized source for best practices, benchmarks and metrics. [20] [21]

CIS had produced a document “The CIS Critical Security Controls”, which gives guidelines on how to prevent the most common and dangerous attacks. Input for the recommendations comes from organizations that are familiar with cyber-attacks, for example U.S law enforcement and top incident response teams including NSA’s Red Team [20]. The guideline was also recently approved and published as a suite of technical reports by European Telecommunications Standards Institute (ETSI) [22].

The document includes 20 critical security controls with multiple, detailed control actions under each clause. Some of the controls are purely technical and some also procedural. The controls are listed according to importance, and the reasoning behind each control is explained. The order is adjusted between versions and new entries added when necessary. The purpose of the list is to distinguish the most important controls from the large amount of security guidance available. The practical experience of the contributing entities ensures that the selected controls are the ones that have been proved to work. The guideline is meant to not only prevent, but also help detecting and restricting ongoing attacks. [23]

The current version of CIS Critical Security Controls is 6.1, released on Aug 31 2016. However, the risk assessment done in this study used version 6.0, since it was the latest at the time of the actual assessment.

3.4 Security Guidance for Critical Areas of Focus in Cloud Computing

Document “Security Guidance for Critical Areas of Focus in Cloud Computing” is published by Cloud Security Alliance (CSA) and unlike the two previous recommendations, it is specific for cloud. The most recent revision 3.0 was released in 2014.

The guideline is roughly divided into two topics: governance and operations. Governance part handles the issues concerning policies and legal issues, whereas operations part is dedicated for the practical implementation and technical aspects. In total the document includes 14 domains [4]:

- cloud computing architectural framework
- governance and enterprise risk management
- legal issues: contracts and electronic discovery
- compliance and audit management
- information management and data security
- interoperability and portability
- traditional security, business continuity and disaster recovery
- data center operations
- incident response
- application security
- encryption and key management

- identity, entitlement and access management
- virtualization
- security as a service.

CSA's guide is primarily directed at those who plan to migrate their services to cloud. Many of the recommendations are not as technically detailed as CIS Critical Security Controls, but they are selected specifically for cloud. Consequently, the controls presented in CSA's document have been picked and balanced with both cloud consumers and providers in mind. The benefit comparing to other controls is that there are no guidelines which would not be suitable for shared technologies.

The guideline emphasizes how the responsibilities between cloud consumer and provider vary between cloud deployment models. In IaaS, the provider can only take care of the layers they offer. Therefore layers starting from the operating system upwards are handled by the consumer. With SaaS the provider offers the whole stack of services, so the responsibilities allocated for the consumer are fewer. However, the division of responsibilities is not always that simple and there are no exact predefined rules. This poses some additional challenges when performing risk assessments. [4]

4 Risk Assessment Process

This chapter describes the theory of risk assessment. The chapter introduces the general framework for the risk assessment process as defined in international standard ISO/IEC 27005:2011, and also the exact methods selected for this study. The methods were chosen based on the data and tools available in the case company's cloud environment and also by referring to the methods used by CSA in their report introduced in Section 2.3.1. There are many threat modelling techniques but STRIDE was selected for the present study since it was also used by CSA's most recent report.

ISO/IEC 27005:2011 standard is part of the ISO27k family and it is complementary to ISO/IEC 27001 introduced in Chapter 3. The standard includes guidelines for information security risk management. The activities described in the standard are risk assessment, risk treatment, risk acceptance, risk communication, risk monitoring and review. The risk assessment process of the framework was used as the basis of this study. The standard does not give specifics on how each step should be performed, so it is up to the organization to decide on the best way to implement them. [24]

The following sections introduce the steps with more details, but a high level description of the risk assessment process is seen in Figure 5.

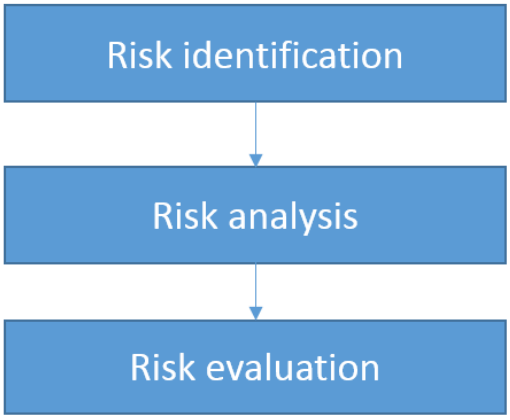

Figure 5. Risk assessment process [24]

Before even entering the risk assessment process, the scope should be agreed on. Defining the scope means deciding what is covered, for example if the assessment is done

to IT infrastructure or to a single application. After that risk identification is performed in order to describe the risks, followed by risk analysis where the probability and impact of the risks are determined. Finally risk evaluation is done where priorities to the risks are assigned and decisions made whether action should be taken or not. These steps are described in more detail in the following sections. Section 4.4 goes beyond the ISO/IEC 27005 framework by introducing the exact methods selected for the study at hand. [24]

4.1 Risk Identification

The first step in the risk assessment process is risk identification. This covers the identification of assets, threats, vulnerabilities, existing controls and consequences. The order of the actions is not dictated in the standard since it depends on the selected methodology. How the identification is done is up to the assessor and organization. [24]

The first item, identification of assets, means defining everything of value to the organization. The assets included do not have to be just software of hardware, but they should be specified within the scope agreed for the risk assessment. After the assets are listed, owner for each is assigned. The owner does not have to be the person who has paid for the asset, but someone who has responsibility for the maintenance, development and security of the asset. [24]

Second task is the identification of threats. The threats can be internal or external and they can occur by accident or on purpose. The methodology for threat modelling is not defined in ISO/IEC 27005:2011, so also in this case it is up to the organization to find a suitable method. However, any previous risk assessments should be consulted if available. There are also different threat catalogues online that can be used as source. Sections 2.2.1 and 2.2.2 introduced some strategies for finding threats. [24]

Fourth item, identification of existing controls, includes listing all the planned and implemented security controls and their statuses. Documentation and discussions with security responsible and users can be referred to. [24]

When vulnerabilities are being identified, the following areas are included [24]:

- organization
- processes and procedures

- management routines
- personnel
- physical environment
- information system configuration
- hardware, software or communications equipment
- dependence on external parties.

A vulnerability can also be the result of ineffective or neglected security control. In this study the security controls introduced in Chapter 3 were used as a guideline.

Identifying the consequences is an important step in order to be able to evaluate the impact of the risks in the next phase of the process. The consequences can be loss of reputation, impacts on effectiveness or actual damage to a system. Cost is not the only factor, instead business effects are also considered. One consequence can affect multiple assets and it can be either temporary or permanent. At this stage the consequences are only listed, not ranked by value. [24]

4.2 Risk Analysis

The second phase of the risk assessment process is risk analysis where the risks are measured. The risks can be scaled quantitatively, qualitatively and semi-quantitatively. In its most basic form risk is calculated as the product of likelihood and impact. Other more complicated formulas can also be used, but in order to get into any conclusion, values need to be known or decided.

In quantitative risk assessment a set of rules and measurable metrics are used to determine the values related to risk. The results can be more precise than with qualitative method, but performing comprehensive quantitative assessment can become very exhausting and require multiple tools. It typically requires historical data from previous attacks or other similar statistics. The downside is that there will be no comparable data for newly discovered attacks, meaning the method may offer false impression of precision. Being able to make the assessment on such a level that it could be repeated, can become rather costly and time consuming and thus outweigh the benefits. The results may also be harder to interpret business-wise if the differences in values are not large. [5][24]

In qualitative assessment the scale is not numerical, but instead values such as low, medium or high are used to describe the risk. Since the scale is not refined, the differences between the levels are not clearly seen. This resulting problem is similar to those in the quantitative where the differences can be too small. With qualitative method two risks that are perceived as medium can actually be quite far from each other. The method also relies on experience instead of metrics and the results may vary depending on who conducts the assessment [5]. However, qualitative analysis can be used as initial assessment before deciding which risks require additional investigation [24]. If the assessment is made by a group of experts and people with understanding on the business needs, the actual assessment may reflect the risk to business more clearly. Communicating the results to management can also be easier with the qualitative method.

The third option, semi-quantitative, can be a good compromise between precision and cost. In this method numerical values are assigned instead of low-high scale. This method communicates risk better than the plain qualitative method, but it too relies on someone assigning the values for each factor used in the assessment. The justification behind each value has to be carefully documented if the assessment needs to be comparable in the future, or if circumstances change and that the outcome needs to be adjusted. [5]

Once the scaling is decided on, it is applied to impact and probability. With impact the consequences of the incident are evaluated. Typically there are technical and business impacts, and it depends on the organization's values, which criteria is used. The value of the asset is taken into consideration, but often the damage cannot be measured only monetarily. The impacts to reputation or lost time and work hours are also a factor, not to mention possible lawful consequences. The impact may have multiple values in cases where, for example, the workhours used for constructing the asset are not substantial, but the value lost if a competitor gets hold of the asset is considerable. The impact may also differ depending if the confidentiality, integrity or availability of the asset is compromised [24].

The second factor analyzed is the probability, or likelihood, of the risk. During the analysis current security controls, vulnerabilities and also motivation of the possible attackers are taken into consideration. In cases where the threat is physical, the location and geography as well as existing physical countermeasures are evaluated. [24]

There is a direct correlation between the likelihood of the security incident and the vulnerabilities of the system, but in order to be exploited, the vulnerabilities have to be discovered. Consequently the skillset and motivation of the attacker becomes a factor as well as opportunities available. That involves understanding whether special access or knowledge is required in order for the exploit to succeed. There are for example attacks that require physical access to the equipment or premises. Furthermore, certain vulnerabilities are easier to discover than others and therefore pose a greater risk. [6]

However, as opposed to technical vulnerabilities, lapses in operational procedures do not even have to be discovered by anyone in advance in order to cause damage. A certain event that normally would not be detrimental for the system can have unexpected consequences even without a malicious intention. For example, an accidental deletion of data from a database, could cause severe damage if backup procedures of the system would not be in place.

Lastly, once the impact and probability are decided on, the risk level is determined. Different kind of matrixes can be formed between impact and probability, or mathematical formulas used to calculate the risk. The result from the task is a list of risks with a value assigned to them [24]. The risk rating is often considered as more of a priority list than that of absolute value. The main purpose is to understand the importance and ranking of different risks.

4.3 Risk Evaluation

After the risk list with values assigned is formed, it is time to evaluate the risks. This step is normally done by the organization and not by the assessor, but it is included here since it is part of the risk assessment process. Evaluation means deciding which actions, if any, should be taken to mitigate the risks.

Factors to be considered include the importance of the asset or process that is affected by the risk and also whether the confidentiality, availability of integrity of the target is important to the organization. There may be legal aspects that affect the decisions as well. [24]

4.4 Selected Methods

As mentioned in the previous section, ISO/IEC 27005:2011 standard does not dictate how threats should be identified. There is no right or wrong way of approaching the task, but there are different strategies. Threat modelling can be run from an asset, attacker or software centric starting point.

In asset centric threat modelling everything starts by listing the things the company considers as an asset. Asset is something of value to the company, for example information or system, but it can also be something attackers see worth stealing. In order for the risk to be realized, the asset actually has to be both. If the information has no value for anyone apart the company, there should not be any reason to go after it. [8]

In this modelling, the threats can be found by understanding ways the assets could be compromised. Consequently, it is vital to have a complete list of the assets, and including only physical items is not enough. Since the assets are often information, understanding the business needs and values can prove important [5]. The downside of focusing on assets is that it can be difficult to make a full list of the assets without understanding the whole system. The analysis can end up focusing only on the things the company sees value in and ignore assets that others may value, thus miss part of the threats.

The second option, attacker centric approach, starts from attackers. This method focuses on who would want to compromise the system and why. The attacker approach can be problematic to implement, since it requires understanding the motivations behind the attack. It can be difficult to make comprehensive list of who might want to negatively impact the system. Therefore this approach is usually not the best starting point [8].

The third option is software centric modelling, where the application is dissected into components and logical data flows drawn between them. Trust boundaries between different areas are drawn and attacks surfaces are better understood. This method requires knowing the software on a level that developers do, but it can give a good overview of the things that may be threatened [8]. In comparison to the asset centric starting point, software centric modelling may actually reveal assets that would have been left out when only covering assets by means of value.

When cloud environments are considered, each analytical approach has its benefits. Software centric may seem like the best option on many cases, but in IaaS there are parts that the cloud provider's own developers may not fully grasp. System usually relies on third party components like the hypervisor and the cloud operating system, which can be difficult to fully model. However, when evaluating a cloud, the assets may not be clear as most of the value comes from the usability and not necessarily from the information itself. Attacker centric starting point on the other hand was excluded since the attacker pool for cloud would have been too large to effectively cover.

Consequently, the approach selected for the present study was software centric. MITRE's CAPEC attack library was studied in order to determine if it would be suitable for finding the threats, but the library was found to be too extensive. It also lacks cloud-specific information, so selecting the correct attacks would have required a longer experience in threat modelling. Therefore, the main threat modelling was decided to be STRIDE, which is also used with software centric modelling. Furthermore, STRIDE was included in the most recent study by Cloud Security Alliance regarding top cloud risks [5]. Even though STRIDE is not designed with cloud environments in mind, other similar studies on cloud have used it [25].

5 Conducting Risk Assessment

This chapter describes how the risk assessment was conducted in the present study. The method for measuring risk in the study was decided to be qualitative. There was not enough data available for making a quantitative analysis, since this was a proof of concept system. The system had not been running before and there was not enough metrics to measure. Vulnerability scans were available but those were the closest to any repeatable measurement available. Furthermore, the scans only covered certain components of the system. The tools that could have been used for this purpose were also limited.

Since this assessment was the first one done to the system, it made sense to make the initial assessment as qualitative. Consecutive risk assessments could then be performed in order to have a closer look at the more serious risks. These assessments could be done as quantitative if possible at that stage.

The actual risk assessment process was divided into five parts. Figure 6 describes the process on a high level.

Figure 6. High level plan of the risk assessment

In the preparation phase the methods for the risk assessment were decided. Two workshops were held with the customer in order to cover all the risks and later to rate them. In the final phase risk assessment report was delivered to the customer. All these phases will be discussed in more details in the following sections.

5.1 Preparation Phase

The goal of the preparation phase was to collect material for the first workshop. Figure 7 shows the tasks in preparation phase.

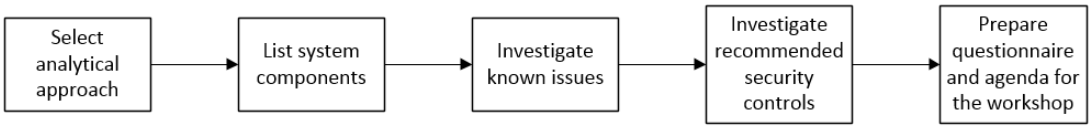

Figure 7. Tasks in preparation phase

The first step was the selection of the analytical approach. As explained in Section 4.4, software centric approach with STRIDE was chosen. After discussing with the cloud provider it was confirmed that that the concept of an asset in this case was too difficult to define explicitly. In addition, starting from the attacker base was seen as a limiting viewpoint, since the PoC system was only available to the company at this phase. Choosing that approach could also make the assessment less feasible considering lessons learned for the production environment.

The second task included consulting the cloud developers in order to get a list of the components used in the system. For example, the software, used to control the creation of the cloud projects and users, was studied in order to understand the system. All other third party product (3PP) components were also listed. Based on the list, the known weaknesses of the software involved were looked at in the third step. Even though MITRE's CAPEC attack library was found to be unpractical in this case, their Common Vulnerabilities and Exposures (CVE) list provided important information on the vulnerabilities of different components. In addition, the web pages related to each function or component were consulted.

During the fourth task, recommended security controls from CIS "Recommendations for Critical Security Controls" [23] and each 3PP software provider were checked. ISO/IEC 27002 standard and CSA's "Security Guidance for Critical Areas of Focus in Cloud Computing" were also consulted for best practices. Supporting questions for the first workshop were selected from these recommendations. Only those that were applicable to the cloud system were listed, but many of the best practices could also be applied to traditional information systems. For example, questions such as "Are the backups stored on a separate location?" or "Is the data encrypted during transfer to the cloud?" were formed.

The purpose of the questions was to guide the workshop to find any possible threats or vulnerabilities in the system, in case dissecting the cloud system and finding threats

based on that would not provide enough information. By referring to the questions, developers could look at different components associated with each function and locate the possible problems. The main idea was still to approach the system via STRIDE method and use the questionnaire as a backup.

5.2 Workshop for Identifying Risks

The goal of the first workshop was to find threats, vulnerabilities and risks related to the system. First the scope was checked in order to confirm what use cases and features should be included and what excluded. Certain functions of the cloud were still in the design phase, and those were directly left out. In the end the scope was limited to the basic functions of IaaS.

Different assets were also defined, but as discovered already in the preparation phase, listing all assets was difficult. Having for example a web server with information databases and a user interface for the customer, would have been more straightforward. In this case the physical machines, networking components, databases and other components were numerous, not to mention the challenges that rise from using shared technologies. Drawing the line between cloud provider and consumer responsibilities was also difficult. Some of the assets first discovered were actually considered to be the consumer's assets when it comes to handling risks, since the provider can influence only on the levels they offer.

An additional problem was the functions that were not yet complete. Some of those used assets that were dedicated for those functions had to be excluded from this assessment, since the use cases were not clear.

After the scope had been defined, the idea of STRIDE was introduced to the developers. Examples of threats related to spoofing, tampering, repudiation, information disclosure, denial of service and elevation of privilege were given to direct the workshop into looking the system from multiple angles.

After the introduction, the developers were asked to draw open the cloud system but it became apparent that visualizing the 3PP parts completely was a challenging tasks. The parts that were developed inside the company and those that were frequently used were easy to understand, but the less used components were left out. Seeing that the cloud

was too complicated and large to be effectively analyzed as a complete system, the questionnaire formed in the preparation phase was taken into use instead.

By following the questionnaire about security controls and requirements a list of the deviations was formed. The supporting list of questions actually revealed components that were not described in the initial attempt of breaking up the cloud into parts. The questionnaire also sprung discussion outside its scope, providing important information.

The deviations from the controls along with other findings from the discussions were written down, and organized after the workshop. Most of the issues discovered were considered to be vulnerabilities. Multiple threats were also found as well as a couple of issues that could be rated directly as risks.

5.3 Forming the Risk Listing

The next phase after the workshop was translating the findings into risks. The case company was not involved in this task. A threat or vulnerability as such does not directly mean a single risk. Many of the vulnerabilities, and in some cases also threats, contributed to multiple different risks. On the other hand, many of them could be grouped under the same risk. As an example, one of the highest risks rated by the previous research in cloud security was shared technology vulnerabilities. If issues with certain software's patch management and cloud tenant separation would have been discovered, both of those would contribute to the shared technology risk. Patch management issues would in addition contribute to the risks related to system vulnerabilities and possibly on the insecure APIs as well.

The final risk listing resulted in for about half of the total issues identified in the workshop. The risks had to be expressed in a level above the threat or vulnerability as there is a difference between those three as discussed in Section 2.2.

Half of the risks discovered could be matched with the risks found from the composite risk listing introduced in Section 2.3.5. The result was expected as the reports included multiple similar findings despite having different sources. If considering the risks from confidentiality, integrity and availability perspectives, all of those aspects were approximately equally represented. Many of the risks were considered to be operational rather than purely technical.

5.4 Workshop for Evaluating Risk

After determining the risks based on the vulnerabilities and threats, a second workshop was held with the case company. In this workshop the business impact and probability of the risks were evaluated together with the developers and system administrators. The case company's input was required since only they have full knowledge about the business impact of each risk. It is viable to suggest rating for the probability based on the findings from the first workshop, but in order to understand the actual impact, in-depth knowledge on the business is needed. On both aspects the opinion of the case company was still considered to be highly influential.

The scale used in the evaluation was negligible, low, medium and high for both probability and impact.

For probability the values were defined as:

- **Negligible:** May occur, but only under exceptional circumstances.
- **Low:** Could occur sometimes.
- **Medium:** Likely to occur.
- **High:** Expected to happen.

For impact the values were:

- **Negligible:** Very low or no impact on the business.
- **Low:** Rather low impact on the business. Business objectives can still be achieved but more resources and/or measures may be needed.
- **Medium:** Clear impact on the business. It is very difficult to achieve the business objectives even with additional resources.
- **High:** Significant impact on the business. It is uncertain that the business can continue.

In this study the system under evaluation was a proof of concept environment, so the business in this case meant the operation of the system. The team evaluated whether the cloud PoC could be used for the purpose it was built.

During the workshop there were three risks that were decided to be left out of the rating. The reason was that the features they concerned were still under development. Those risks would have to be rated after the implementation was closer to being completed.

After determining the impact and probability, the risk rating was calculated based on the risk matrix shown in Figure 8:

IMPACT	RISK LEVELS				
High	H	H	H	H	
Medium	M	M	H	H	
Low	L	L	M	H	
Negligible	I	L	M	M	
	Negligible	Low	Medium	High	PROBABILITY

Figure 8. Risk rating matrix

The risk levels were divided into informational, low, medium and high. If the impact of the threat was high, the risk level was always high regardless of the probability. This reflects the importance of the risk assessment for the business function. If the probability of the threat was high, the risk was also high unless the impacts were negligible. The risks where the total rating was high were considered unacceptable and the rest acceptable. Unacceptable in this case means that actions should be taken immediately.

The risk levels with required action plans were as follows:

- **Informational:** Minor risk or observation. Should be followed in long term
- **Low:** Limited risk. Should be followed and actions considered in long term
- **Medium:** Poses a risk to operations and requires actions also short term
- **High:** Serious risk to operations. Requires immediate action

5.5 Risk Assessment Report

After the risks had been ranked according to severity, a report was written to the customer. The report included description of each risk and its probability and impact. The

related threats, vulnerabilities and existing controls were also listed along with the recommendations.

Before delivering the report to the customer, the risk ratings were also revised by senior security consultants in the company. Some of the risks were adjusted compared to the initial workshop with the customer. All of the risk levels that were altered were raised to a higher risk level. The difference was mostly caused by the different viewpoint on PoC system. Customer's original rating was based on the fact that the target was not a production system, but the security consultants wanted the risk levels to give better feedback for the upcoming production environment.

6 Discussion

In this chapter the results of the study are discussed. The goal of this thesis was to investigate how risk assessment is conducted and also evaluate the suitability of the selected methods for cloud environment. The exact findings of the risk assessment were not the main focus, instead it was the procedure of making the assessment for a cloud environment.

In general, performing a qualitative risk assessment is not an exact science. It is affected by the competence and experience of the assessor. The product team providing input for the risks assessment is also a vital part of a successful evaluation. The larger the system, the harder it is to effectively cover everything. If the target is an application or a single server, it may be easier to assess the system with purely technical methods. Vulnerability scans and penetration testing as well as historical data on previous attacks, or at least security monitoring, would have been helpful in making the results repeatable and more reliable. Since the cloud system was still being developed those were not available at this time. For this study, the method of relying on workshop data as well as system component best practices was the correct choice.

The original plan of using STRIDE threat modelling as the basis of the workshop was not optimal. STRIDE would have suited better for assessing an application or maybe a smaller system. Using STRIDE for the whole cloud infrastructure proved to be difficult and most of the input was received by going through the questionnaire related to the recommendations by CSA, SANS Institute and ISO/IEC 27002. In all, the problem was not with the selected threat model, but limitations of modelling in general. Trying to split the cloud system into manageable pieces for the workshop was not practical. Groundwork could have been made this way, but for the discussions it was not the best choice. More emphasis could have been placed into building the questionnaire had this been known. In this study the questionnaire was meant as supporting material instead of a foundation.

However, having a workshop for discovering risks was found useful. Relevant information was discovered during the workshop and the team responded openly to questions concerning the recommended best practices and controls. It could be argued whether the same questionnaire could have been delivered on paper or electronically without face-to-face discussion, but the information collected during the workshop exceeded the

scope of the questions. Having multiple people looking at the recommendations and thinking of deviations produced more value than one person answering the questions alone. For many of the topics the initial answer from one person started a discussion that led to a result that was different from what was originally thought. The probable reason for that is that a single person rarely holds a full understanding of the system.

One thing to consider when having a workshop for discovering risks is who should participate. In this assessment all the participants were technical personnel, including the project responsible. Some of the topics related to operational procedures could have benefited from having higher management present. However, that could have affected how openly the rest of the team responds to questions and the dynamics of the discussion. A good solution would have been having people from development, administration and operational sides, preferably ones that are closely involved in the daily operation of the system. In the internal PoC phase, these departments were not yet formed, so the participants were limited.

When it comes to the workshop for rating risks, it could have been valuable to include people from the managerial side. The business impacts would have been best rated by those that have an overall understanding of the economics. However, in this phase the business value could not be measured with costs or profits, instead only functionality of the system could be evaluated.

In the rating workshop the competence of the assessor is essential. With more experience, the assessor is able to effectively question the probabilities and impacts suggested in the product team. Especially probability can be justified with technical knowledge and previous experiences from other systems. In this study the initial rating relied too heavily on input from the workshop team. An assessor with a longer experience in risk assessments would have been able to contribute more to the workshop.

The actual findings of the risk assessment were not the main focus here, however the results are discussed on a general level in order to understand how the selected methods may have affected the outcome. Figure 9 shows how the risk severities were divided.

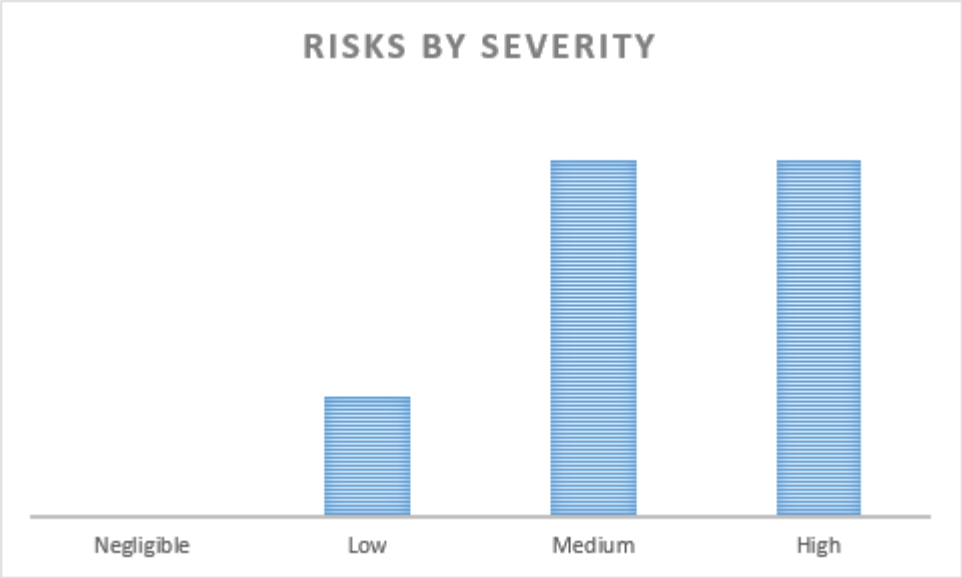

Figure 9. Number of risks by severity

The values were removed from the y-axis due to confidentiality reasons, but the chart gives as overview of the risks discovered. The rating shows the results after the severities were revised by senior security consultants.

75% of the high and medium risks discovered were more operational than technical. It can be assumed that this was affected by the method used in the workshop. The questionnaire related to best practices included more recommendations on the operational than technical level. Therefore it is logical that the risks discovered reflect that. Furthermore, since the assessment was done in the PoC phase and the rating adjusted to correlate with severities of a production environment, it is expected that the findings include more high ranked risks related to operational procedures than what would be discovered in the production phase. The reasoning is that in the PoC phase the system and operations around it are not fully matured. There are certain processes that are only implemented as going to production.

Another factor affecting the results is the lack of complete vulnerability analysis of the components used in the system. Even though MITRE's CVE database and vendor web pages were consulted, it does not equal to a full vulnerability analysis. Therefore some of the technical threats were most likely not discovered or at least their probability not evaluated to full extent.

Finally, the stage where the assessment was done should be discussed. Making the risk assessment during internal PoC phase had its benefits and disadvantages. Not all the features were fully implemented and some were still in the design phase. The evaluation at this stage cannot be said to be comprehensive, but the assessment gave valuable feedback on how to improve the cloud environment before going to production. Even large changes to the system were still possible when the risk assessment was done this early. The basic functions of the cloud environment were in place and improvements on them could be made with moderate effort compared to having to rethink designs close to launch. Consequently, making the assessment in the internal PoC phase was valuable, but new assessment should be made as more features are completed.

7 Conclusions and Recommendations

This chapter goes through the conclusions from the study and also presents recommendations for future risk assessments. The different methods for finding threats are compared and improvement plan for the next risk assessment is introduced.

7.1 Conclusions on Threat Modelling

Three main sources were used for when finding threats. This section compares those methods. STRIDE, CAPEC attack library and a checklist of recommendations on security controls were examined during this study. The last one is not strictly speaking a threat modelling method, but it was used to discover threats, vulnerabilities and risks in this assessment. The comparison between the methodologies is shown in Table 6.

Table 6. Comparison of methods

Method	Advantages	Disadvantages
STRIDE	Can discover emergent, non-defined threats	Requires good description on functions and flows in the system
	Very flexible	High level model
	Does not constrain	
	Excellent for small/well-defined applications	
CAPEC	Very detailed	Hard to get into
	Extensive list	Time consuming
	Real world attacks, not theoretical	Restricts creativity
		Requires good understanding of the system
		May not discover emergent issues
Security control checklists	Good for large systems	Dependant on discussion
	Real world experiences as basis, not just theory	May not discover emergent issues
	Suitable for beginner	Not a proper threat model

It should be noted that the advantages and disadvantages in Table 6 may change depending on the system and experience of the person doing the modelling. For cloud environment, where the environment is very complicated and drawing open the data flows and functions was difficult, STRIDE suffered the most. The benefits of STRIDE are

undeniable and it allows discovering new threats better than the other methods. If the assessment was not done to IaaS system and all the functions were listed meticulously, STRIDE would have been the best method. It does not limit threats to those already known and allows for great flexibility. Like checklists, STRIDE could also be considered to be dependent on discussion with the developers, but experienced modeler would probably still get comprehensive list even from less discussion. If modeling a more constrained system like a database or a web server, STRIDE would be the recommended method.

Unlike STRIDE, CAPEC attack library focuses on attacks instead of security properties, so comparing those two directly is not straightforward. CAPEC library was extremely detailed, but the amount of attacks listed was overwhelming for a person who has not done threat modelling before. If done by one person, going through the whole library would have taken unreasonably long. Since the system was a cloud environment with partly unclear use cases, CAPEC was also difficult to apply. Finding relevant threats from CAPEC requires good technical understanding of the functions involved. For more experienced assessor, the library may also act as a constraint and not encourage to think outside the norm. Similarly to STRIDE, CAPEC could be efficient when examining a system with clearer limits. For such system, combining STRIDE with CAPEC would probably be a good idea assuming the modeler is already familiar with some of the attacks in CAPEC and the time invested in using it would be justified.

The third option, security control checklists from standards and other recommendations, proved to be the best option when assessing IaaS cloud. As mentioned earlier, the discussions with developers were extremely valuable and resulted in discovering multiple threats. Since the discussion relied on controls for existing attacks, it is likely that new attacks would not be found by using the checklists. However, considering the information available from the system and the experience of the assessor, the method was suitable. The biggest downside is that the risks discovered would have been a lot less if the discussions were not so open.

Had it been possible to define the use cases as simple functions and draw those open, STRIDE could have been used for those while the IaaS infrastructure was evaluated with the help of the check lists.

7.2 Recommendations for Future Risk Assessments

Risk management is a continuous process and therefore risk assessments should also be regularly updated. As new features are being implemented, the system should be re-assessed. The same applies if the system will go to production.

If going to production, it would be recommended to collect additional data from vulnerability scans, penetration testing and security monitoring of the system. As a result there would be more reliable data for the risk assessment which would not be as dependent on the assessor or the team taking part in the assessment. During vulnerability analysis the system is examined with technical methods in order to discover weaknesses. Penetration testing on the other hand takes it one step further and tries to exploit the weaknesses and gain access to the system. Penetration testing is therefore more intrusive than vulnerability analysis and requires different kind of expertise. Often this is performed by ethical hackers. Both of these actions are separated from the actual risk assessment and should be done prior to starting the assessment.

Based on experiences from this study, an improved plan for the risk assessment is shown in Figure 10.

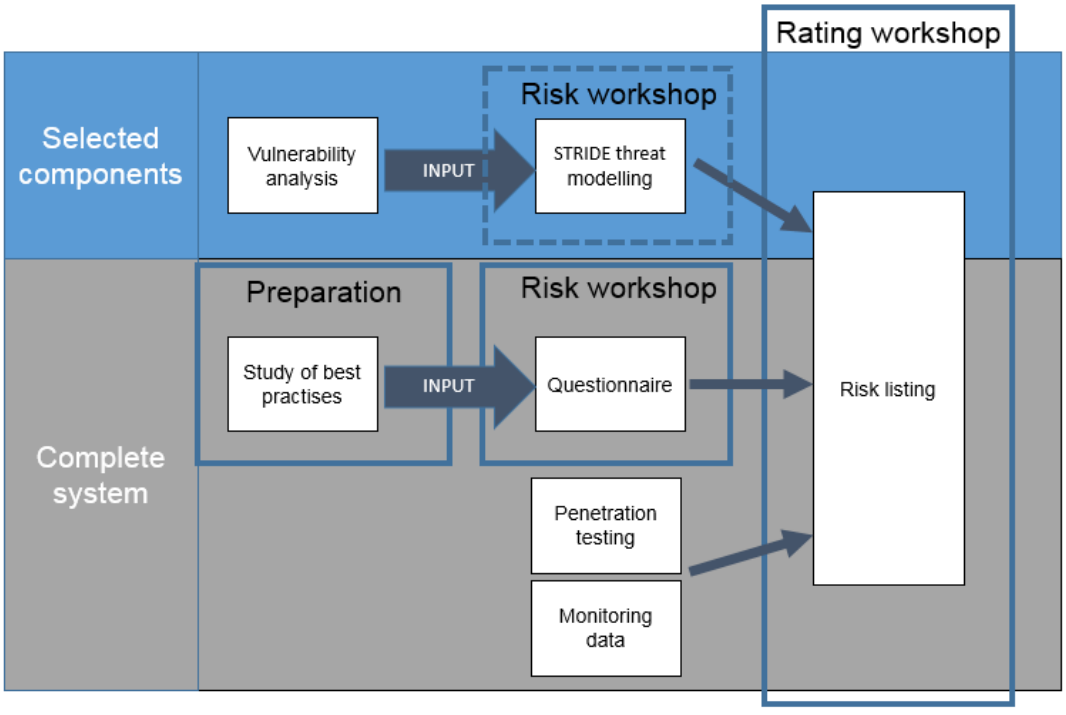

Figure 10. Improved plan for risk assessment

As displayed in Figure 10, it would be recommended to gather data from vulnerability analysis of selected components before going into the risk assessment. As discovered during this study, splitting the whole system into components may not be feasible with a large cloud system, but there are certain functions that are more clearly separated and could be analyzed individually outside the assessment. For example user portal, identity management, storage API and cloud management API are such functions.

After the data is collected from the vulnerability analysis, the selected components could be examined via STRIDE threat model. This can be done as a workshop with the developers assuming the functions are built in-house. Optionally, the findings of the vulnerability analysis can be directly included in the risk listing.

As for the complete system, the method of using a questionnaire with best practices in a workshop was found effective in this study. The same approach could be complemented with penetration testing of the system and data collected from security monitoring and vulnerability scans. Information from all of these sources could then be included in the risk listing.

Finally a workshop for rating the risks can be held. By combining the different data sources, the listing would not rely as heavily on the assessor. The rating should still be discussed with the cloud provider, since they do have the best understanding of the business functions. However, it would be best if the assessor is able to do a preliminary rating which would give a starting point for the discussions.

References

1 Verizon. 2016 Data Breach Investigations Report (DBIR) [online]. 2016. URL: http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/. Accessed 13 Jun 2016.

2 National Institute of Standards and Technology. Special Publication 800-145 - The NIST Definition of Cloud Computing. [online]. Sep 2011. URL: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf. Accessed 24 May 2016.

3 National Institute of Standards and Technology. Special Publication 800-146 - Cloud Computing Synopsis and Recommendations [online]. 21 May 2012. URL: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf. Accessed 18 Mar 2016.

4 Cloud Security Alliance. Security Guidance for Critical Areas of Focus in Cloud Computing v3.0 [online]. 14 Nov 2014. URL: https://downloads.cloudsecurityalliance.org/assets/research/security-guidance/csaguide.v3.0.pdf. Accessed 27 May 2016.

5 National Institute of Standards and Technology. Special Publication 800-30 - Guide for Conducting Risk Assessments [online]. Sep 2012. URL: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf. Accessed 25 May 2016.

6 Open Web Application Security Project. Risk Rating Methodology [online]. 3 Sep 2015. URL: https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology. Accessed 25 May 2016.

7 International Standard. ISO/IEC 27005:2012 - Information technology - Security techniques - Information security management systems - Overview and vocabulary. Second edition. 1 Dec 2011.

8 Adam Shostack. Threat Modeling: Designing for Security. Feb 2014. Wiley. p. 29-104.

9 Open Web Application Security Project. Threat Risk Modeling [online]. Updated 11 Apr 2016. URL: https://www.owasp.org/index.php/Threat_Risk_Modeling. Accessed 6 Jun 2016.

10 MITRE. Common Attack Pattern Enumeration and Classification – About CAPEC [online]. Updated 30 Jun 2014. URL: https://capec.mitre.org/about/index.html. Accessed 1 Oct 2016.

11 MITRE. Common Vulnerabilities and Exposures - About CVE [online]. Updated 13 Sep 2016. URL: https://cve.mitre.org/about/. Accessed 11 Oct 2016.

12 Cloud Security Alliance. About: Cloud Security Alliance [online]. Not dated. URL: https://cloudsecurityalliance.org/about/. Accessed 21 May 2016.

13 Cloud Security Alliance. Cloud Security Alliance Releases 'The Treacherous Twelve' Cloud Computing Top Threats in 2016 [online]. 26 Feb 2016. URL: https://cloudsecurityalliance.org/media/news/cloud-security-alliance-releases-the-treacherous-twelve-cloud-computing-top-threats-in-2016/. Accessed 21 May 2016.

14 The European Union Agency for Network and Information Security. About ENISA [online]. Not dated. URL: https://www.enisa.europa.eu/about-enisa. Accessed 22 May 2016.

15 The European Union Agency for Network and Information Security. Cloud Computing Benefits, risks and recommendations for information security. Nov 2009. URL: https://resilience.enisa.europa.eu/cloud-security-and-resilience/publications/cloud-computing-benefits-risks-and-recommendations-for-information-security. Accessed 6 May 2016.

16 Open Web Application Security Project. About The Open Web Application Security Project [online]. Updated 22 May 2016. URL: https://www.owasp.org/index.php/About_OWASP. Accessed 21 May 2016.

17 Open Web Application Security Project. OWASP Cloud – 10 Project [online]. Updated 23 Jan 2014. URL: https://www.owasp.org/index.php/Category:OWASP_Cloud_%E2%80%90_10_Project. Accessed 21 May 2016.

18 International Standard. ISO/IEC 27001:2013 Information technology – Security techniques - Information security management systems – Requirements. Second Edition. 1 Oct 2013.

19 IsecT Ltd. ISO/IEC 27002 Code of Practice [online]. 2015. URL: http://www.iso27001security.com/html/27002.html. Accessed 14 Sep 2016.

20 SANS Institute. SANS Institute - CIS Critical Security Controls [online]. Not dated. URL: https://www.sans.org/critical-security-controls. Accessed 20 Sep 2016.

21 Center for Internet Security. Center for Internet Security – About [online]. 2016. URL: https://www.cisecurity.org/about/. Accessed 20 Sep 2016.

22 Infosecurity magazine. ETSI Incorporates CIS Controls [online]. 19 Aug 2016. URL: http://www.infosecurity-magazine.com/news/etsi-incorporates-cis-controls/. Accessed 20 Sep 2016.

23 Center for Internet Security. The CIS Critical Security Controls for Effective Cyber Defense, version 6.0 [online]. 15 October 2015. URL: https://www.cisecurity.org/critical-controls/. Accessed 1 Jan 2016.

24 International Standard. ISO/IEC 27005:2013 Information technology – Security techniques - Information security risk management. Second Edition. 1 Jun 2011.

25 Katerina Lourida, Antonis Mouhtaropoulos, Alex Vakaloudis. Assessing Database and Network Threats in Traditional and Cloud Computing. International Journal of Cyber-Security and Digital Forensics Vol.2(3) pp.1-17. 1 May 2015.

ENISA - Cloud computing risks

Policy and organizational risks:

- R.1: Lock-in
- R.2: Loss of governance
- R.3: Compliance challenges
- R.4: Loss of business reputation due to co-tenant activities
- R.5: Cloud service termination or failure
- R.6: Cloud provider acquisition
- R.7: Supply chain failure

Technical risks:

- R.8: Resource exhaustion (under or over provisioning)
- R.9: Isolation failure
- R.10: Cloud provider malicious insider – abuse of high privilege roles
- R.11: Management interface compromise (manipulation, availability of infrastructure)
- R.12: Intercepting data in transit
- R.13: Data-leakage on up/download, intra-cloud
- R.14: Insecure of ineffective deletion of data
- R.15: Distributed denial of service (DDoS)
- R.16: Economic denial of service (EDoS)
- R.17: Loss of encryption keys
- R.18: Undertaking malicious scans or probes
- R.19: Compromise service engine
- R.20: Conflicts between customer hardening procedures and cloud environment

Legal risks:

- R.21: Subpoena and e-discovery
- R.22: Risks from changes of jurisdiction

- R.23: Data protection risks
- R.24: Licensing risks

Risks not specific to cloud:

- R.25: Network breaks
- R.26: Network management (ie. network congestion, mis-connection, non-optimal use)
- R.27:Modifying network traffic
- R.28: Privilege escalation
- R.29: Social engineering attacks (ie. impersonation)
- R.30: Loss or compromise of operational logs
- R.31: Loss or compromise of security logs (manipulation of forensics investigation)
- R.32: Backups lost, stolen
- R.33: Unauthorized access to premises(including physical access to machines and other facilities)
- R.34: Theft of computer equipment
- R.35: Natural disasters