NTNU

Norwegian University of
Science and Technology

Security in Cloud Computing

A Security Assessment of Cloud Computing Providers for an Online Receipt Storage

Mats Andreassen
Kåre Marius Blakstad

Master of Science in Computer Science
Submission date: June 2010
Supervisor: Lillian Røstad, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Problem Description

We will survey some current cloud computing vendors and compare them to find patterns in how their feature sets are evolving. The start-up firm dSafe intends to exploit the promises of cloud computing in order to launch their business idea with only marginal hardware and licensing costs. We must define the criteria for how dSafe's application can be sufficiently secure in the cloud as well as how dSafe can get there.

Assignment given: 14. January 2010
Supervisor: Lillian Røstad, IDI

Abstract

Considerations with regards to security issues and demands must be addressed before migrating an application into a cloud computing environment. Different vendors, Microsoft Azure, Amazon Web Services and Google AppEngine, provide different capabilities and solutions to the individual areas of concern presented by each application. Through a case study of an online receipt storage application from the company dSafe, a basis is formed for the evaluation. The three cloud computing vendors are assessed with regards to a security assessment framework provided by the Cloud Security Alliance and the application of this on the case study. Finally, the study is concluded with a set of general recommendations and the recommendation of a cloud vendor. This is based on a number of security aspects related to the case study's existence in the cloud. With dSafe's high demands of data locality, integrity and security, Google AppEngine is discarded as an option due to the lack of focus on business related applications, whilst Microsoft Azure is the recommended cloud vendor – closely followed by Amazon Web Services – due to its suitable technical solutions with regards to existing implementation, risk mitigation capabilities and audit results.

Preface

This report is the result of our Master Thesis work during the spring of 2010.

Throughout the work we have gained much insight into both the positive and negative aspects of cloud computing as well as knowledge of the capabilities of three of the major cloud computing vendors. We foresee that we will benefit greatly from this experience during the course of our careers.

We would like to thank our supervisor, Lillian Røstad, for her guidance and Daro Navaratnam, as well as his team, for allowing us to use dSafe in our case study and their assistance during said study.

Kåre Blakstad and Mats Andreassen

Trondheim, June 2010.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Cloud Computing has long been a *buzzword* within the field of computer science. As with most buzzwords there are numerous definitions, but most agree that Cloud Computing entails putting applications and data in the hands of others. Unless you run your applications and keep your data on your own hardware on-premises, you leave the security of said assets in the hands of others. In a world where applications are riddled with weaknesses and vulnerabilities, can you really trust another company to keep your assets?

> *When a computer is within your network, you can protect it with other security systems such as firewalls and IDSs. You can build a resilient system that works even if those vendors you have to trust may not be as trustworthy as you like. With any outsourcing model, whether it be cloud computing or something else, you can't. You have to trust your outsourcer completely. You not only have to trust the outsourcer's security, but its reliability, its availability, and its business continuity.*
> – Bruce Schneier*[1]*

There are several providers[1] available and their platforms differ substantially in their flexibility. How much of the infrastructure the customer is able to access can directly affect the security of their own services as well as that of others. This difference can have critical consequences: Research has been described in which meticulous analysis of a cloud's hardware in idealised circumstances allow attackers to log the keystrokes of users on separate virtual machines [2]. Some of the cloud platforms are therefore polished surfaces with little access to the underlying infrastructure and others allow customers to delve deeper.

This report will detail these differences and their consequences for security. By comparing the platforms we intend to, from a security standpoint, comment on the maturity of a selected set of current cloud computing vendors. In addition, we will be performing a case study on the start-up business dSafe. dSafe intends to create a

[1] In this report we will use *vendor* and *provider* interchangeably.

information system with the entire Norwegian people as potential users and deploy it off-premises, in the care of a cloud vendor. dSafe has tentatively chosen the Windows Azure platform. Our intention is to find out if dSafe's planned system can be deployed securely to the cloud. There are several aspects to this problem, perhaps the most important of which are the technical and legal. In both, utilising cloud technology has clear ramifications. To complicate matters, as we will see in the beginning of Chapter 2, there is only a somewhat consensus on what cloud technology actually is.

In the remainder of this chapter, we present our research methodology.

1.1 Research Questions

The research questions establish the scope of the study.

Q1: Can information systems be as secure in the cloud as they would be in an on-premises environment?

Has the cloud matured to the point where it is possible to create reasonably secure cloud systems? That is, at least as secure as non-cloud systems. The core issues of confidentiality, integrity and availability become extra important whenever outsourcing data servers due to the inclusion of an external data provider.

Q2: How can the planned services of dSafe be sufficiently secured in the cloud?

We will analyse dSafe's business plan as well as design documents to find out how dSafe can deploy its intended services sufficiently secure on the platform of their chosen vendor as well as those of other vendors. What we deem to be *sufficiently secure* for dSafe will also be explored as a part of the aforementioned analysis.

The next section will elaborate on how we intend to answer these questions.

1.2 Approach

Our answering of the research questions consists of three steps:

1. Comparative study of cloud vendors (Chapter 2).
2. Understanding the case (Chapter 3).
3. Case study of an online receipt archive (Chapter 4).

Even though some quantitative data exists in most of the cases (e.g. how many applications that run on the different platforms), these do not let us extrapolate any useful conclusions regarding security. Most of the sources we base our discussion on are either research papers, technology papers, our own experience with the platforms or other users' experiences. In addition we consult the publicly available information and documentation on the respective vendors' websites.

1.2.1 Comparative Study

In order to answer our research questions we first need to understand the challenges that await in the cloud as well as familiarise ourselves with cloud vendors. We intend to proceed in the following steps:

- **Introduce cloud computing:** First we present what we deem to be essential knowledge and terminology within the field of cloud computing.
- **Introduce cloud vendors:** In Section 2.3 we consider vendors and present the capabilities of the ones we deem relevant. The number of vendors studied should not exceed three, in order to limit the scope of the study. These vendors are selected by applying the following criteria:
  - The vendor should be a large market participant within the cloud, and other areas based on data centers, to ensure future availability of services.
  - The vendor's platform can accommodate dSafe's planned system.
  - The vendors should be different in nature, to ensure that not only vendors, but also different cloud approaches, are explored.
  - Sufficient information must be present on the vendor, in order to do a qualified reasoning about the vendor's properties relevant to the application.
- **Investigate security aspects:** In Section 2.4 we identify security aspects and investigate the different vendors' capabilities within each aspect.
- **Summary:** Finally, we intend to find some patterns in the vendors' capabilities and intentions.

1.2.2 Understanding The Case

In understanding the case, we have had four meetings with dSafe in which we discussed their business plans and design documents. During the course of our writing, we keep email correspondence ongoing to be able to keep up to date with dSafe's progression. The documents and the other information acquired form the basis for our dSafe presentation in Chapter 3.

1.2.3 Case Study

In this study we utilise a framework laid out by the Cloud Security Alliance: Security Guidance for Critical Areas of Focus in Cloud Computing [3], specifically its second major revision. The Cloud Security Alliance (CSA)'s mission statement is:

> *To promote the use of best practices for providing security assurance within Cloud Computing, and provide education on the uses of Cloud Computing to help secure all other forms of computing.* – CSA

A fairly new security endeavour (conceived late 2008) for a fairly new field. The Alliance has quickly gained a great many affiliate and corporate members. This includes affiliates such as The Open Web Application Security Project (OWASP) and Jericho Forum, while the corporations include Microsoft, Google, Dell, HP, Symantec, VeriSign and VMWare [2]. The paper itself is co-authored by a cohort of security experts from the different member organisations divided into into worker groups. These worker groups are responsible for the guidance within the 13 different domains as defined by the master paper [3].

The first domain, *Cloud Architecture* details the field of Cloud Technology and introduces the terminology. This domain, as well as the following six *governance domains* and six *operational domains* are intended to exhaustively cover any and all security aspects that companies should consider when contemplating deploying information systems in the cloud. We have found that some of the domains have overlapping aspects and so we do not follow the guidance too stringently as we make our recommendations to dSafe.

In order to answer question two, we intend to follow the following steps:

1. **Identify and classify assets:** In Section 4.1 we determine any and all assets dSafe has based on the information presented in Chapter 3, both current and future. After all assets have been identified we determine the risk associated with six scenarios. These scenarios are related to information disclosure and outside manipulation and are detailed in Section 4.1.2.
2. **Deployment Model Acceptance:** In Section 4.2 we determine what cloud model is acceptable for the identified assets.
3. **Analyse dSafe in the context of domains:** In Section 4.3, within the context of the twelve domains, we make recommendations for how dSafe should proceed in order to mitigate the risks we have uncovered.
4. **Risk Mitigation:** Finally, in Section 4.4 we summarise what vendors provide mitigations within each domain to visualise the of each platform.

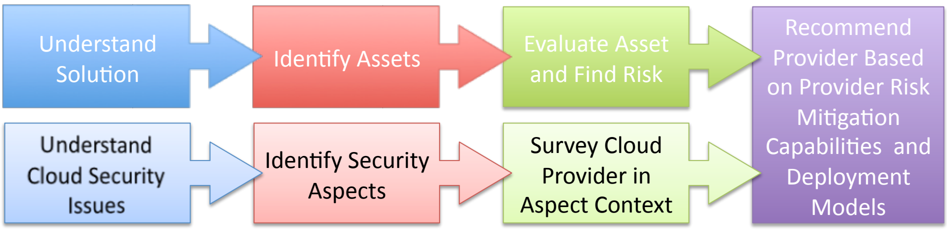
Understand Solution
Identify Assets
Evaluate Asset and Find Risk
Recommend Provider Based on Provider Risk Mitigation Capabilities and Deployment Models
Understand Cloud Security Issues
Identify Security Aspects
Survey Cloud Provider in Aspect Context

Figure 1.1: The process of recommending a provider based on the case's application assets, the providers' security aspect capabilities and adherence to deployment models.

[2] See `http://www.cloudsecurityalliance.org/Membership.html`

[3] See `http://www.cloudsecurityalliance.org/guidance.html`

1.3 Report Outline

The report is divided into five parts.

- **Introduction:** This chapter, which introduces the problem description as well as our methodology for solving the problem.
- **Cloud Computing:** The chapter explores what cloud computing is and what it means for software security. It will include an analysis of what security challenges the cloud poses and what mechanisms the different providers have implemented to meet said challenges.
- **Case Overview:** The chapter explores this newly started firm. This includes the firm's business plan, their vision, and an analysis of their security challenges.
- **Case Study:** The chapter contains our analysis of their business idea, *an online receipt archive*, and the viability of placing their services in the cloud while maintaining the security of their clients.
- **Conclusions:** In the final chapter we intend to answer the research questions posed in this chapter, as well as describe what further work that needs to be done.

Chapter 2

Cloud Computing

For decades, companies have been buying and maintaining their own computational infrastructure for hosting Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), web pages and business web portals. This infrastructure has been scaled to handle the largest amount of traffic realistic to the business. For online shopping this means that the infrastructure must handle the Christmas shopping season, for weather reporting portals this means the amount of traffic usually present after a hurricane and for universities this means the enrollment period, where all student are required to register for classes. The rest of the time, the system usually works at a sub-optimal utilisation, meaning that the organisation has idle computational resources. If enterprises with non-overlapping high-demand time periods were to share the computational power and further utilise this more uniformly, one would approach a more effective usage of the resources. The centralised maintenance and operation of the computers would also entail a more cost-effective operational model, due to the larger degree of specialisation such a computer farm organisation would allow compared to the average in-house IT operator.

But what about the different platforms and technologies needed by the different users of this centralised computer center scheme? One clearly needs different environments running on these servers in order to meet the need of every customer and to guarantee some sandboxing and protection of data and processing done for each user. This is solved through virtualisation. In short, virtualisation is *one computer acting like many*. Some virtualisation software run on the hardware while the operating systems and user-specific software runs in *instances* on top.

In order for these services to run seamlessly, one is dependent on some kind of fault tolerance and redundancy. If all users were running on top of a single point of failure, the Service Level Agreement (SLA) would not seem very appealing for anyone to accept. Therefore, one is dependent on running multiple computers with seamless backup switching, replacement of fault hardware, adding more computational power if needed – all transparent to the user. To make this happen, the computer farm providers use a technique called *grid computing* which is *making many computers act like one logical unit.*

Essentially, *cloud computing* is a combination of *virtualisation* and *grid computing.* In essence, this scheme allows the end user to not care about the underlying infrastructural architecture. The virtualisation part allows the user to have an exclusive environment to work on and the grid computing part allows the system to be as scalable and fault tolerant as expected for a dedicated system. While this description of cloud computing is simplified and focused on the infrastructural architecture, other definitions exist:

> *Dynamic provisioning of IT capabilities over the network.*
> – Henrik Lund-Hanssen, Master Architect, Accenture
>
> *Cloud computing is a way of computing, via the Internet, that broadly shares computer resources instead of having a local personal computer handle specific applications.*
> – Wikipedia.org
>
> *What is cloud computing all about? Amazon has coined the word "elasticity" which gives a good idea about the key features: you can scale your infrastructure on demand within minutes or even seconds, instead of days or weeks, thereby avoiding under-utilization (idle servers) and over-utilization (blue screen) of in-house resources. With monitoring and increasing automation of resource provisioning we might one day wake up in a world where we don't have to care about scaling our Web applications because they can do it alone.*
> – Markus Klems, Student
>
> *People are coming to grips with Virtualization and how it reshapes IT, creates service and software based models, and in many ways changes a lot of the physical layer we are used to. Clouds will be the next transformation over the next several years, building off of the software models that virtualization enabled.*
> – Douglas Gourlay, VP of Marketing, Arista Networks
>
> *Cloud computing is a model for enabling 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.*
> – Peter Mell and Tim Grance, National Institute of Standards and Technology
>
> *Cloud computing is about gracefully losing control while maintaining accountability even if the operational responsibility falls upon one or more third parties.* – Cloud Security Alliance [3]

These definitions differ somewhat, but the concepts of resource utilisation and how virtualisation allows the user to focus on the services repeat themselves as the most important advantages of cloud computing throughout these definitions. The final definition ignores these aspects, instead focusing specifically on the security aspect.

This chapter consists of the following sections:

- **Service Models:** This section presents and establishes the three different service models from the literature: Software-as-a-Service, Platform-as-a-Service and Infrastructure-as-a-Service.
- **Deployment Models:** This section presents the five different cloud deployment models: public, on- and off-premises private, community and hybrid.
- **The Cloud Vendors:** Here we present some cloud vendors and their different components. These are the providers we discuss throughout the the rest of this report.
- **Cloud Security Aspects:** The cloud infers many challenges for the company wishing to put their applications and data in it. There are many different aspects to consider from the security perspective: E.g. security audits, certifications, security incidents or resource pooling.
- **Summary:** Finally, we compare the vendors to look for patterns between them.

2.1 Service Models

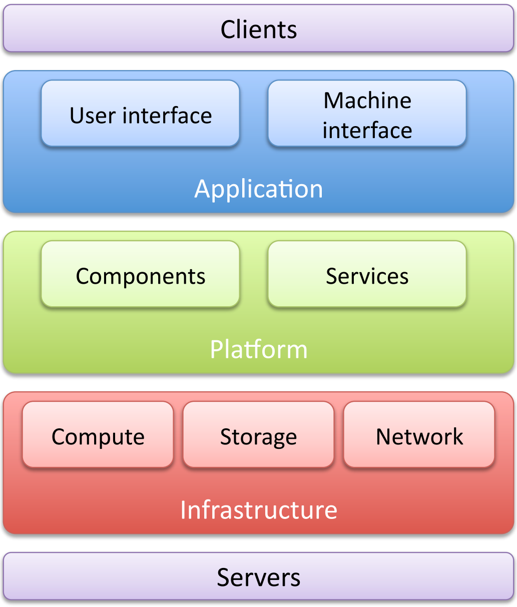

Figure 2.1: Overview of the cloud computing layers

As the array of definitions has shown us, cloud computing seems to be a rather abstract concept and is usually concretised into more tangible categories. If a service is running on *the cloud* it can usually be called cloud computing as long as it runs on a specialised computer system with virtual environments, or spaces, for each user. This encompasses not only applications, such as Google Mail or Docs, but also the virtualisation of operating systems, rendering the user more freedom compared to the application schema. As such, the cloud computing services is usually concretised into these categories based on the taxonomy defined by Youseff et al.[4], also shown in Figure 2.1:

- Application layer: Software-as-a-Service (SaaS)
- Platform layer: Platform-as-a-Service (PaaS)
- Infrastructure layer: Infrastructure-as-a-Service (IaaS)

Regardless of which definition of Cloud Computing one accepts it might sound similar to outsourcing. Hence a clarification is useful:

> *Cloud Computing in general can be distinguished from traditional outsourcing in three ways: the time of service (on-demand and intermittent), the anonymity of identity of the service provider(s) and anonymity of the location of the server(s) involved.* – Cloud Security Alliance[3]

As we move from SaaS through PaaS to IaaS this difference becomes more evident. The following quick treatise on the different categories will explain further. Figure 2.2 attempts to categorise Google AppEngine, Amazon Web Services and Windows Azure in the cloud computing service layer model.

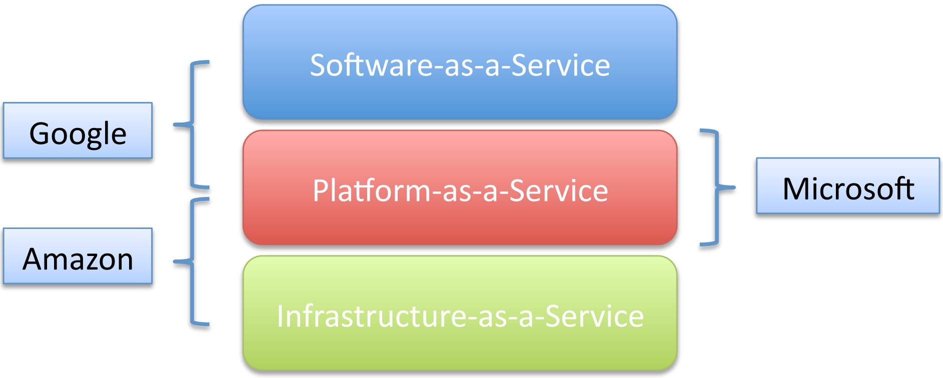

Figure 2.2: Categorisation of Google AppEngine, Amazon Web Services and Windows Azure According to the Cloud Service Models.

2.1.1 Software-as-a-Service

These services run on the cloud application layer. A service provider license application access based on demand. The service is usually hosted in an environment controlled fully by the service provider, eliminating the need for the user to install and maintain the application. Examples are Google Mail, a mail application only accessible by web, and SalesForce.com, a business software company, running all applications on a hosted off-site purchased through subscriptions to the service. This scheme allows end users to neglect the cost of hardware, licenses and the maintenance of the infrastructure. According to [5], the SaaS segment is the only segment proven to be successful as a business model.

2.1.2 Platform-as-a-Service

In this layer, a platform or a solution stack is provided that exists on top of the cloud infrastructure which is transparent to the user. The platform service segment facilitates the cloud applications which the user may deploy in this environment. Examples are Windows Azure, the Google AppEngine and Amazon Web Services. If one needs an environment on which to deploy and host custom application, the PaaS offers more dynamic functionality compared to the rigid and application specific environment of the SaaS. Consequently, the customer is more directly responsible for the reliability and security of the service provided to the end user.

2.1.3 Infrastructure-as-a-Service

The infrastructure is the most fundamental service segment. The suppliers of this service provide the storage and processing capabilities needed to run the abovementioned services. The IaaS may include managed hosting, development platforms or otherwise clean virtual machine instances reachable through the remote vendor. In this segment as well, Amazon is a big market operator, but also other well-known companies, such as IBM, are present with managed hosting services. Further, not much effort is usually needed in order to find a local supplier of remote virtual machine environments. In the IaaS category, the customer is even more responsible for the reliability and security of the services they provide.

2.2 Deployment Models

To differentiate between cloud implementations considerably different in nature, a set of deployment models have been defined – similar to how the service models (see Section 2.1) have been defined. We will give a short run-down on the different models here based on the National Institute of Standards and Technology (NIST) definition [6, 3].

2.2.1 Public cloud

This cloud infrastructure is accessible by the general public or a cluster of organisations. This system is hosted, managed and owned by an organisation selling

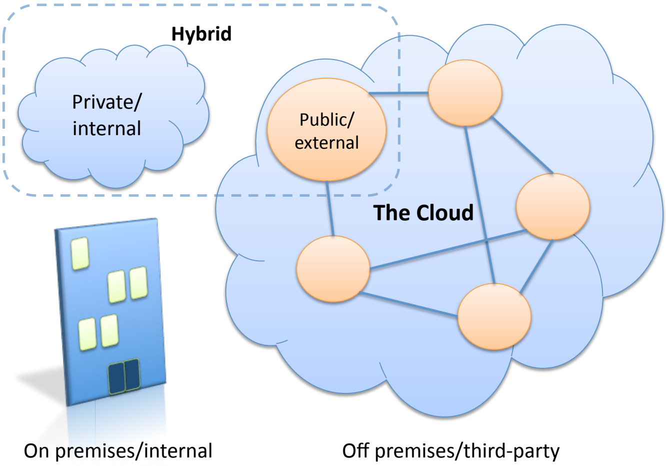

Figure 2.3: Overview of the cloud computing deployment models

cloud services. This means that a multitude of people and corporations put their applications and storage into the same cloud infrastructure. Entities looking for cloud capabilities should consider this fact immensely before deciding on a cloud provider.

2.2.2 Private cloud

This cloud infrastructure is operated solely for one organisation. There are two types of private clouds:

- **The private internal cloud:** The organisation acquires the necessary hardware and maintains it for itself.
- **The private external cloud:** The organisation pays a cloud provider to provide this as a service.

2.2.3 Community cloud

This cloud infrastructure is operated for multiple organisations with a set of shared concerns – typically cooperative, mission- or domain specific concerns. As with the private cloud, this environment may be managed or hosted by the organisations, or any third party vendors.

2.2.4 Hybrid cloud

This cloud infrastructure is a combination of two or more of the above models. The clouds composing a hybrid cloud remain unique entities and are bound together by technology allowing communication between the clouds. Applications of this technology can include load-balancing and application portability, which makes the hybrid useful in different contexts, including high-load environments or when processing redundancy is needed.

2.3 The Cloud Vendors

In this section, we will explore some providers of cloud technology. As briefly stated in the introduction (Section 1.2) these specific providers are chosen based on some criteria. All three selected providers fit all these criteria based on a qualified reasoning on todays standing in the cloud computing market. In summary the vendors chosen for this comparative study are:

- **Windows Azure:** Fairly new, it graduated from beta status January 1, 2010.
- **Amazon Web Services:** The oldest platform in the study, launched in July of 2002.
- **Google AppEngine:** The middle child, launched April 7, 2008.

2.3.1 Windows Azure

Windows Azure is a platform that provides a Windows-based environment for running applications and storing data in Microsoft data centers. In short, Windows Azure is a an operating system available as an online service. The platform runs on a large number of machines accessible though the internet. On top of this environment, the computation and storage services exist as own layers. A thorough overview on the platform can be found in David Chappell's white paper on the subject from 2009 [7]. The platform consists of three core components:

- **Computing:** A complex solution built on Microsoft's HyperV virtualisation technology, Azure is able to supply a flexible number of virtual machines for running the actual applications.
- **Storage:** Azure additionally supplies different types of storage devices to suit the developers' needs.
- **Fabric:** Finally, there is the *AppFabric*. This fabric is the environment in one of Microsoft's data centers that controls all the hardware necessary to supply the complex services of the first two components.

The computing services are two-fold. Both are virtual machine instances, but only the *web role* instances are meant to be visible to end users of the applications. The *worker role* is meant to "do the heavy lifting". By increasing the number of instances and controlling their access patterns through a load balancer, Azure provides intelligent scalability for developers. To enable communication between these instances, Azure provides a queueing service. All components of the Azure platform are designed as REST-ful web services that can be accessed through HTTP/HTTPS.

Microsoft has created an elaborate web site through which developers can deploy, maintain and supervise their applications. Finally, as signified by the Azure slogans “I know a place that’s different, but familiar” and “I know a place where I can code in my language” developers can run any applications that runs on the Windows platform. That is, even though the web role instances come bundled with Microsoft’s IIS web server, nothing stops developers from installing Apache Tomcat and running their enterprise Java applications.

2.3.2 Amazon Web Services

Amazon Web Services (AWS) is a collection of infrastructure services. Some of the most important of these are:

- **Elastic Compute Cloud (EC2):** A service for uploading, managing and hosting a virtual machine instance – Amazon Machine Image (AMI) – and providing a scalable capacity for this instance.
- **Simple Storage Service (S3):** A web service for storing and retrieving potentially large amounts of data from the underlying Amazon infrastructure.
- **CloudFront:** A service for distributing content from the S3-services to the end user. While supplying dynamic content through some service (EC2), the CloudFront may deliver static content such as images.
- **SimpleDB:** A web service for running queries on structured data and provides the core functionality of a database.
- **Simple Queue Service (SQS):** A hosted queue for holding messages awaiting to be sent between components in a distributed application.

Most of the AWS, and all of the above, are not services made directly available to the end user, but offered as functionality for the developers to use. The services are available over HyperText Transfer Protocol (HTTP). The pricing of the Amazon Web Services is based on how much it is used. Amazon offers a scalable model, with the possibility to scale the computational power as needed for the applications hosted on the EC2 service.

2.3.3 Google AppEngine

Google AppEngine (GAE) is a platform for developing and hosting user-developed software on infrastructure owned and managed by Google. Until now, Google has not targeted enterprise customers specifically. On May 19th of 2010, Google held it’s annual I/O convention. Among the many things presented was the very relevant *Google AppEngine for Business* keynote. For Google to be able to target enterprises they have created an SLA. However, the current SLA is merely a draft and the enterprise features are still in a beta state.

- **Application Hosting:** The platform supports an array of technologies with main focus on Java Virtual Machine (JVM)-running languages and Python, and some support of popular web frameworks for these. The platform is hosted on a virtualised grid system, making Google able to scale the applications transparently for the user.

- **Datastore:** Google's datastore is their approach to create an alternative to the relational database in the cloud. The data is partitioned and distributed to achieve redundancy and performance enhancements. Due to the challenges of running queries against multiple related tables in an SQL database, Google has implemented a datastore syntax called GQL. This new syntax puts some restrictions on the selections, e.g. joins are not supported due to the inefficiencies of multi-machine queries [8]. Select statements in GQL can therefore be performed on only one table.
- **Hosted Structured Query Language (SQL):** Still in GAE's future. This is Google's challenger to become a contender in providing reliable cloud functionality for businesses.
- **Blob Storage:** During the 2010 I/O convention, Google announced *Google Storage for Developers* [9]. This service is intended as a direct challenge to Amazon's Simple Storage Service (S3). Data put into this storage is replicated to enable redundancy and scalability and is available through a Representational State Transfer (REST)ful interface. At time of writing the service is part of Google Labs (a beta program) and only available to a limited number of North American developers.

2.4 Cloud Security Aspects

> *Cloud computing ought to be called swamp computing and we don't even know what the alligators are yet.* – Adrian Seccombe

The vendors and their technologies are presented in Chapter 2. After establishing which cloud providers should be a part of the study we then choose different security aspects and surveyed the vendors in each aspect. In determining which aspects to investigate we tried to be exhaustive and we consulted an array of sources. The major ones are:

- Gartner's *Seven cloud-computing security risks* [10].
- Jericho Forum's *Commandments* [11].
- The Cloud Security Alliance's *Security Guidance for Critical Areas of Focus in Cloud Computing* [3].

The top two sources contain subsets of the bottom one. The latter source is very thorough and we have been unable to locate any other security aspects not covered within. By looking at the vendors through the goggles of each security aspect we intend to discover whether there are any major differences between how they approach security on the whole.

These groups frequently tout their members, and others – large providers of cloud computing, such as Amazon, Google and Microsoft, both directly and through media. An increased focus from these companies will surely benefit the end user when it comes to security. A focus on security and the adoption of taxonomies and standards set by respected specialist organisations does not only gain the user of the service, but also the cloud providers themselves. According to Schneier[1], cloud computing is no different from other computing other than the trust you have in the provider. In order for the large providers to gain this trust, it is critical that advice originating from a specialist organisation, such as CSA is taken into account – since these organisations are typically comprised of both large providers and large consumers of cloud computing.

In this section we investigate the capabilities and stances of the different vendors from Section 2.3. The aspects we have chosen are the following:

- **Service Level Agreements and Compliance Features:** The different providers have different SLAs that define what amount of downtime is deemed acceptable. What are the main differences between vendors and do any of them even make guarantees? All customers must accept these agreements if they wish to deal with these providers.
- **Authentication Services:** These services are essential in providing integrity, confidentiality and access control to both customers and the users of their applications, whether in the cloud or not.
- **Audit and Certifications:** What guarantees do customers have as to the manner in which cloud vendors process the data entrusted in their care? For a company to choose to be certified is a strong, public signal that they take their responsibilities seriously.
- **Cloud Tech Support:** Outsourcing data center operations means that the customer does not have full control over their information system. When security incidents such as attempted break-ins occur; who is contacted?
- **Incident Response and Logging:** The information system with human or machine interfaces will be constantly subjected to both illegal and legal access. Since the hardware and software that runs the information system in the cloud it is vital that customers are notified of security incidents in the cloud and are able to respond properly without being hindered by lack of physical access.
- **On-demand Self-service:** This aspect of cloud computing can cause complications for the security policies of the customer. Several customers on the same hardware may have adverse effects on their ability to access log information and some vendors might not let customers perform penetration testing on their production systems.
- **Broad Network Access:** This aspect is about how cloud systems are designed to allow the consuming of services from numerous platforms and devices.
- **Resource Pooling:** Outsourcing data center operations means leaving data center operations in the hands of a third party vendor. The vendors pool

resources in data centres across the globe and this might have legal consequences for customers. Does the customer have any control as to where the data goes?

- **Internal Access Control:** Do the vendors have internal controls in place for authorising administrative personnel to access the customers data? The customer needs assurances to this effect.
- **Virtualisation:** For multi-tenant cloud services to be provided at infrastructure, platform or application level, the use of virtualisation is inevitable for the provider to remain cost efficient. It is the multi-tenant part that can adversely affect the security of customers' applications. Does the virtualisation technology compartmentalise sufficiently?

2.4.1 Service Level Agreements and Compliance Features

The different providers have different SLAs that define what amount of downtime is deemed acceptable. What are the main differences between vendors and do any of them even make guarantees?

All customers must accept these agreements if they wish to deal with these providers. The amount of time that services are available divided by the total time it was supposed to be available becomes the percentage known as uptime. This uptime is often referred to as a "number of nines". For instance, four nines means 99.99 %. This equates to about 53 minutes out of 365 days and is certainly very impressive. We shall now see what guarantees, if any, the providers make regarding service levels. We will also take a look at what features the different providers have implemented to achieve their SLA-stated goals. These intentions regarding uptime might seem impressive but the SLAs makes it clear that the service credits are the only compensation the customer is eligible for. Note that each SLA describes these *Service Credits*. The credits are percentages of the usage fees that are reimbursed when certain criteria regarding uptime are reached. The downtime is usually calculated as an accumulation of five minute periods a service is unavailable. We will say otherwise when that is not the case. Any downtime occurring for less than five minutes is not counted however many such outages occur.

Windows Azure

Microsoft has four different SLAs which each define uptime and compensation for failure to satisfy these limits:

- **Azure Compute:** For both *monthly connectivity downtime* and *monthly role instance uptime*, the system is designed to reach at least 99.99 % uptime. If it drops below this, a 10 % service credit is given. This service credit further increases to 25 % when it falls below 99.95 % [12].
- **Azure Storage:** Service credits are 10 % below 99.9 % and 25 % below 99 %. The method for calculating this monthly uptime is 100 % −

average error rate . The average error rate is based on a set of operations and their acceptable completion time. Further details can be found in the Azure Storage SLA [13].

- **SQL Azure:** A period of five minutes is determined as “unavailable” if all the customer’s connections fail or take more than 30 seconds to conclude. The number of unavailable periods is counted and the monthly availability service level is determined in a straight forward manner. The limits are also here 99.9 % and 99 % with the corresponding service credits [14].
- **AppFabric:** The calculation is equivalent to the one described for SQL Azure. The connectivity in question here is the one between “either the Access Control Service or the Access Control Service Management Service endpoints and Microsoft’s internet gateway” [15].

Microsoft has put in place an impressive set of measures to provide these availability and reliability goals [7]:

- **Load Balancing:** In Windows Azure, all resources such as virtual machines (web and worker roles) and storage nodes are protected by load balancers. This is done to prevent single instances from being overworked.
- **Redundancy:**
  - All the different types of storage available (Blobs, Tables and Queues) in Windows Azure are all backed up with *triple redundancy* as of this writing [16].
  - *Fault Domains* are something Microsoft has implemented transparently for the user. In effect, the hardware in Microsoft’s data centers are divided into these domains. When services scales up the AppFabric makes sure to spread the instances evenly between the fault domains in the data center. In this way they avoid that when certain hardware fails not every instance of a customer’s service fails.
  - *Upgrade Domains* are somewhat related to Fault Domains in that worker and web role instances in upgrade domain one can be upgraded while the ones in domain two are upgraded after a) the instances in domain one are fully upgraded and b) new traffic is directed to domain one. This makes for close to painless upgrades of instances in Microsoft’s cloud.

Amazon Web Services

Amazon has two somewhat different SLAs for their two types of services [17, 18]:

- **AWS S3 Storage:** When the uptime percentage drops below 99.99 % the client is eligible for a 10 % service credit on his monthly billing. When it drops below 99.95 %, the discount is increased to 25 %.
- **AWS EC2:** When the uptime percentage drops below 99.9 % uptime, the client is eligible for a 10 % service credit on his yearly billing.

To achieve these uptime levels Amazon claims that "AWS will use commercially reasonable efforts". There are two kinds of features for achieving this: implemented by Amazon and implementable by the customer [19].

- **Load Balancing:** *Elastic Load Balancing* is a transparent feature implemented by Amazon. It hides all your identical instances behind an alias, taking care of balancing the load between the instances [19].
- **Name Service:** When new instances are created they are automatically given both a public and a private IP address. The customer may exchange this public address with an *Elastic IP* address instead. Elastic addresses are a kind of dynamic IP address that overcomes that problem of Domain Name System (DNS) propagation latency on the public internet.
- **Redundancy:**
  - The EC2 cloud is divided into *availability zones* that comprise *regions*. By placing storage and computing instances in several regions across availability zones the customer is able to influence and increase the reliability of his services since these zones are maintained as completely independent of each other. Amazon provides Application Programming Interfaces (APIs) that enables what they call Auto Scaling. As the name implies this allows the customer to program in rules for how the number of EC2 and S3 storages are duplicated to cope with increasing and decreasing load. The *CloudWatch* feature gathers metrics to enable this rule programming.
  - The *Amazon Elastic Block Store* is a persistent storage volume that can be attached to EC2 instances. These volumes are stored with a redundant copy inside availability zones. When computers fail for different reasons the storage volumes connected to it may end up with less than intact integrity. Amazon's block store is designed so that the integrity of volume is sustained even when EC2 instances fail. This enables new instances to immediately connect to the storage volume without delay.
- **Protection:** Amazon has implemented several protections against Denial of Service (DoS) attacks [20]: *SYN cookies*[1] and connection limiting. In addition Amazon maintains internal bandwidth that is greater than the Internet Service Provider (ISP)-given bandwidth given to EC2 and S3 instances to at least make sure that internal communication can continue unhindered.

Google AppEngine

For its enterprise customers Google has now provided a tentative SLA [21]. Similarily to the other providers it aims to provide *Service Credit Refunds* when the service level drops below predefined boundaries. When the monthly service level drops below 99,9 % the customer receives a 10 % discount, below 99 % yields 25 %, below 95 % yields 50 %. What sets Google's scheme apart from the others is that

[1] See `http://en.wikipedia.org/wiki/SYN_cookies`

when the service level drops below 90 % the discount becomes 100 %. The same kind of accumulation of five-minute periods of downtime as the other providers do is applicable here.

Google provides automatic scaling and load balancing but few details have been published [22]. Among the available APIs that are published for Java and Python a simple blacklist filter for DoS protection is described. The recent publication of their *Google AppEngine for Business* program indicates that they will open up further in the future as enterprises start pressuring them for more information and features.

2.4.2 Authentication Services

These services are essential in providing integrity, confidentiality and access control to both customers and the users of the customers' applications, whether in the cloud or not. Before this age of cloud computing the development of such features have rested completely with the developers. We will see if any of the cloud vendors have prepared facilities for assisting their customers with this. In building general infrastructures the vendors become uniquely suited to affect the security of customers' applications that run on it.

Windows Azure

Microsoft has implemented something simply called Access Control (AC). This is based on on a specification called *WS-Federation*[2] which was co-developed by IBM, Microsoft, Novell and VeriSign and details how federated identity can be implemented across web services. It details mechanisms for allowing information brokering on identities and authentication [23]. AC implements this *claims-based identity model.* This model is also detailed in a presentation from the Jericho Forum [24], but Microsoft's implementation is detailed in [25]. A simplification of the model follows: In such a model the user of the service, i.e. the store, authenticates with an authenticating service that supplies a token for use with the real services provided. The token is a set of attributes (called *claims*) concerning an identity, for instance a username, an e-mail address or a set of privileges. In this manner one avoids transferring the password every time a receipt upload occurs. Instead the token is passed whenever operations are performed until the token expires. The tokens are designed to include digital signatures to enable quick and easy validation. This can then be used as a kind of "single sign-on" service that can benefit developers greatly. They can now use this service across their services without having to implement disparate authentication into every web service. This, arguably, scales better than the alternative. In addition, it allows developers to authorise partners, e.g. banks, to issue such tokens as well, to ease cross-service integration.

[2] See `http://www.ibm.com/developerworks/library/specification/ws-fed/`

In 2009, Microsoft passed their first Security Assertion Markup Language (SAML) 2.0 interoperability tests, and will most likely in the future be interoperable with other technologies supporting this eXtended Markup Language (XML)-format. SAML includes an eGovernment profile. To date, a number of governments have employed this standard. A federated identity is not a requirement for this service, but when sharing identity with widely used, trustworthy services, this would most likely increase the usability of the service and encourage further use.

As Microsoft provides full platform virtualisation in Azure, any technology running on Microsoft Windows Server 2008 runs in the instance. Therefore, any third party authentication schemes compatible with the Microsoft platform will work.

Amazon Web Services

Amazon has authentication against the various supported services in the AWS stack. Since the Amazon cloud is mostly an infrastructure cloud, and any platform can be hosted, it leaves authentication outside the AWS environment up to the developer – making most authentication implementations possible – but with the added development overhead.

Google AppEngine

The Google App Engine only have native support of the OAuth[3] authentication protocol. OAuth is now implemented throughout the Google Data API, including Google Docs, Calendar and YouTube. Although proved working and reliable even for financial transactions (through projects such as OpenTransact and projects implementing this library[4]) OAuth has not yet been adopted by the larger financial institutions [26]. OAuth, and the similar OpenID has yet to be adopted by security intensive institutions in Norway due to the lack of trivial security measures, such as two-factor authentication. There exist solutions mapping OpenID's to more personal identities[5]), such as BankID, but OpenID is not deemed secure enough for financial and semi-sensitive data. Google App Engine also has an API for the SAML 2.0 authentication protocol, described in the Windows Azure subsection.

Since AppEngine supports an array of technologies, any authentication tactic available for implementation through these available technologies is possible.

2.4.3 Audit and Certifications

The main goal of the cloud providers, the data center owners, is to make money. This will always be the top priority. One can not assume that the provider will have the best interest of its customers as its top priority. Normally, the provider will

[3] See http://oauth.net/

[4] See http://www.opentransact.org/

[5] Provided as an unofficial product from Norwegian company Signicat (http://signicat.com/)

have little direct access to the customers' data beyond any contracted agreements between the two parties, but in some cases the provider might take liberties with the customer data, which would not very likely fall in the customers good grace. The customers must take responsibility for asking the right questions regarding control routines and obtaining knowledge about the potential areas for security vulnerabilities.

To obtain information on the provider's internal controls, regardless of the certifications obtained by the cloud provider, it is essential to obtain some commitment or permission for the customer to conduct external third-party audits. Data center internal controls can be tailored to pass some certification, and the details of the audit are seldom released. The customer should not have to place any trust in the professional relationship between an internally hired auditor and the cloud provider.

Due to the strict access and disclosed procedures, a cloud provider will most likely not accept audit requests from the regular customer – this would not only breach standing protocols regarding data center security, but would also break the promises made to the other customers with regards to the security and integrity of their data. If a large and business critical customer were to request such a review, the provider would most likely be more lenient on the security protocols. A thorough – and expensive – audit performed by a large company could also benefit both the data center provider and the smaller customers. If an audit-program were available for all customers, this would be highly preferable. Certifications such as SAS70 and ISO 27001 have been created so that independent parties can perform audits and certifications, thereby enabling customers to make industry recognised assumptions in this field.

Both Amazon and Microsoft offer the geolocalisation setting as described in Section 2.4.8 but Amazon does not publish the exact addresses of their data centers. They do this to provide a sort of "security by obscurity" layer on top of all their other security measures but it also effectively prevents even dSafe from inspecting the on-site security. Google does not even allow for specifying location, so the customer would have no idea which data center to inspect.

In 2005 International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC) completed a standard called *ISO/IEC 27001:2005 - Information technology – Security techniques – Information security management systems – Requirements*. We will refer to it simply as ISO 27001. Since this standard formally specifies a set of requirements to be satisfied, it becomes possible to audit and certify organisations. The certification process consists of three steps, the first of which is an informal review of the organisation's Information Security Management System [27]. This must consist of an information security policy, a Statement of Applicability and a Risk Treatment Plan. The Statement of Applicability defines security control measures, the why and where they are needed [6].

[6] See `http://iso-17799.safemode.org/index.php?page=Statement_of_Applicability`

The Risk Treatment Plan is the result of thorough risk assessments [7]. ISO 27001 recommends the *Plan Do Check Act*-cycle to develop these items. In this iterative approach you start with charting risks and find security controls, actually implement them, evaluate their effectiveness and finally make the necessary changes to optimise security. Step 2 is an actual independent compliance audit to find out whether the developed Information Security Management System (ISMS) satisfies the requirements of ISO 27001. This would then result in an actual certification. Step 3 consists of subsequent follow-up reviews to ensure that the organisation keeps adhering to the process, possibly leading to decertification. The CSA is "issuing an industry call to action to align cloud providers behind the ISO/IEC 27001 certification". Furthermore, when companies become ISO 27001 certified it is possible to gain insight into the process, thus learning about the information security policies of the vendor, by submitting a written request [28].

The SAS 70 certification is a certification created in the USA by the American Institute of Certified Public Accountants (AICPA). No organisation based outside the United States may issue certifications. The CSA underlines that a SAS 70 Type II audit (the most thorough one) only serves to create a snapshot in time of how the implementation of security protocols corresponds to the documentation.

Windows Azure

> *Independent, third-party validation OSSC's (Online Service Security and Compliance) include Microsoft's cloud infrastructure achieving both SAS 70 Type I and Type II attestations (...)* – Charlie McNerney, GM, Business & Risk Management, Microsoft Global Foundation Services

Microsoft is also ISO 27001 certified [29].

Amazon Web Services

A quote from the AWS security whitepaper states the following after listing compliances and certifications given under internal third-party audits:

> *(...) These certifications provide outside affirmation that AWS has established adequate internal controls and that those controls are operating efficiently. AWS will continue efforts to obtain the strictest of industry certifications in order to verify its commitment to provide a secure, world-class cloud computing environment.*

It seems that Amazon regards these compliances as an affirmation of AWS internal controls. Implicitly this states that these internal audits is sufficient to confirm the quality of these controls. We have been unable to locate any public statement regarding Amazon and ISO 27001 certification.

[7] See `http://iso-17799.safemode.org/index.php?page=risk_treatment_plan`

Google AppEngine

Google's SAS70 Type II certification is intended to ensure that the following protocols and controls are in place, as stated by Google:

- **Logical security:** Provide assurance that only authorized personnel may access Google Apps production systems.
- **Privacy:** Controls provide assurance that procedures and policies regarding the privacy of the customer data in relation to Google Apps.
- **Data center physical security:** Controls ensuring that Google Apps are hosted in a protected enveironment.
- **Incident management and availability:** Ensures that Google provides redundancy to increase availability and that incidents are reported and treated correctly.
- **Change management:** Controls providing reasonable assurance that testing and code review is done prior to any changes in the production environment.
- **Organisation and administration:** Controls that provide assurance that management provides the infrastructure and mechanisms to follow up on initiatives within Google that impact AppEngine.

We have been unable to locate any public statement regarding Google and ISO 27001 certification.

2.4.4 Cloud Tech Support

Essentially outsourcing data center operations to a third party entails that customers must understand their options when things go awry. The cloud vendors are corporate entities who's prime motivation have been pleasing their shareholders. While this might not have changed the vendors seem to recognise the cloud customers' need for support. We will here survey the the support capabilities of the vendors.

Windows Azure

Microsoft's support capabilities are somewhat difficult to summarise but they offer mostly the same features as the other providers. Developers can choose different types of subscriptions that dictates their eligibility for different levels of support. All professional support is handled through the "Microsoft Online Services Customer Portal". The features available are online reporting, call-back feature, 1:1 support, 24/7/356 availability and Azure-specific forums within the Microsoft Developer Network for community support [30].

Amazon Web Services

AWS also offers three different levels of support very similar to Google's offering [31].

- **Free:** *Community Forum*, *Service Health Dashboard*, and *Frequently Asked Questionss (FAQs).*
- **Premium - Silver:** Business day support according to Pacific Standard Time, 1:1 support via web based ticketing system.
- **Premium - Gold:** 24/7/365 availability, 1:1 support via telephone, 1 hour guaranteed response time on urgent issues.

The premium features pertain to all the infrastructure services such as EC2 and S3.

Google AppEngine

Google offers three kinds of support [32]:

- So far, developers have only had access to **Community Support** for which Google makes no guarantees. This is basically a forum where developers discuss issues and Google responds to questions from time to time.
- A new model is **Operational Support**. Google guarantees 1:1 support, online forms for submitting issues, a targeted 1 hour response time for critical issues, as well as 24/7 availability for critical issues. This type of support is confined to system outages, billing & quotas and account management.
- The final model is **Premium Developer Support**. The same 1:1 support is guaranteed, a private support portal provided, targeted response times of 8 hours, and the availability is confined to what is defined as business hours in Pacific Standard Time. This type of support is for technical guidance issues.

2.4.5 Incident Response and Logging

The information system with human or machine interfaces will be constantly subjected to both illegal and legal access. Instrumentation is by many considered to be a form of art. Fundamentally it comes down to trading some performance for the availability of system traces. The art lies in what one does with these traces. These system traces or logs can be used for several purposes. We will see that cloud providers are very good at providing performance metrics. What is more interesting for the purpose of this study are security metrics. These security metrics allow for incident analysis and prevention schemes. Such systems are often referred to as Security Information and Event Management systems. The implementation of such a system will allow companies to:

- Define and detect security events.
- Create rulesets for dealing with said events.

- Structure and aggregate different types of logs for auditing.

Windows Azure

The Windows Azure Software Development Kit (SDK) contains several APIs. It allows the programmatic management and retrieval of a lot of information: Windows Azure logs (on by default), IIS 7.0 log, Windows Diagnostic infrastructure logs (on by default), Failed Request logs, Windows Event logs, Performance counters, Crash dumps and Custom error logs [8]. All of these items of diagnostic information can be retrieved by any application through a REST-ful API, allowing monitoring to occur from both outside and inside the cloud. The format of the data returned from these requests is a well documented XML structure.

Amazon Web Services

Amazon CloudWatch enables the corresponding functionality in AWS. The AWS Management Console is for instance able to show graphs of CPU, Disk and Network metrics but all this data is also available as a REST-ful web service. This data is aggregated and persisted for two weeks before this temporary storage is recycled. The format of Amazon's returned data is very similar to Microsoft's, being a human readable XML structure.

Google AppEngine

Google offers an *AppStats* API for both their Java and Python platforms. This API lets the application's owner view request logs and performance metrics:

> *Appstats retains statistics for the most recent 1,000 requests (approximately). The data includes summary records, about 200 bytes each, and detail records, which can be up to 100 KB each.* – Google

For both language platforms Google has created way of enabling what they refer to as the *management console* for applications. On this console, the metrics are used to generate graphs and statistics. The API exposes considerably fewer metrics than Microsoft's or Amazon's. This seems logical since Google does not expose any infrastructural control to its users and does not even publish their infrastructure metrics on a regular basis.

2.4.6 On-demand Self-service

This is the ability to provide auto-scaling of the compute and storage cloud depending on instance load. For a manually sized instance, this would mean API for watching the load and scale the instance accordingly. For a dynamic instance, this

[8] See `http://msdn.microsoft.com/en-us/library/ee758705(v=MSDN.10).aspx`

would be the ability to closely review the load of the instance and to set an upper limit of resources assigned to the instance.

Windows Azure

Windows Azure provides the *Service Management REST API*[9] for programmatic scaling and management of instances. Combined with one or more of the Azure logging API (see Section 2.4.5 for further details), one can determine the needed number of instances at any given time. This is not available as an automatic feature, since it would entail a complex cost model, giving Microsoft complete control over the costs related to traffic, storage and computational time. Correct use of these API will provide both rapid elasticity, depending on the algorithm used on the data interpreted by the measurement service, also provided though the Microsoft API.

Amazon Web Services

Amazon EC2 offer the same capabilities as Azure through the *Amazon CloudWatch* services. This set of APIs provide both services for logging – measured services – (see Section 2.4.5 for further details) and for managing the instances and scaling the capabilities. Amazon does not provide this as a built-in feature, because of the cost-scaling implications this would imply.

Google AppEngine

In the AppEngine price model, you pay for what you use, not for what you reserve *and use*, as the two above providers. Also, the other providers emulate individual machines that you can control, Google does not attempt to virtualise individual machines. This leaves the AppEngine customer with less flexibility, less control and less options, but makes it easier for the application programmer, rendering efficient resource utilisation almost transparent. Therefore, services for measuring capability resources and scaling these are not available for the AppEngine. Since you have all the compute space in the world available, but only pay for the resources you use, there is no reason to worry until the data traffic to you application really takes off and the amount paid to Google starts to look substantial. An upper limit of cost can be defined through the administrative interface to prevent out-of-budget scaling.

2.4.7 Broad Network Access

Broad network access means *to what degree is computational and storage capabilities available from different platforms.* As more and more internet services use XML-based, standardised protocols – available over HTTP – the services become more

[9] Service Management REST API: `http://msdn.microsoft.com/en-us/library/ee460799.aspx`

and more accessible from most devices. This section will focus on HTTP-based protocols.

Windows Azure

The Windows Azure cloud is strictly public, all resources are meant to be accessed through web service interfaces (so-called RESTful services). This makes it very easy to transfer data in or out of this cloud infrastructure.

Amazon Web Services

For AWS the same applies. Resources are meant to be interconnected with the use of provided REST and Simple Object Access Protocol (SOAP) interfaces. Amazon also offers Virtual Private Cloud; a private off-premises cloud alternative that offer the same services as the public that can be accessed through IPSec. This offering is primarily meant for organisations wanting to put parts or all of its intra-organisational applications in the cloud and not for public services.

Google AppEngine

Contrary to the storage solutions of AWS and Windows Azure GAE's current storage services are not connectable from outside the cloud, only the actual applications in the cloud may do this programmatically. This is not a show-stopper for dSafe since their storage units were never intended to be accessed outside their applications. Google does not make recommendations for RESTful services.

2.4.8 Resource Pooling

Resource pooling means grouping together resources, e.g. storage, processing, memory, network bandwidth and virtual machines, in order to maximise advantages and profits and/or minimise risks. A cloud provider may specify the location of these resources to some extent, at a higher lever of abstraction (e.g. country, city), but at the same time availability and data center redundancy promises are made. So where is your data being processed, stored and sent? We will take a look at each provider to identify if any of the providers can guarantee the location of the customer instance.

Windows Azure

When deploying services on Windows Azure, a geo-location option is given in the deployment UI. One can assign the application to a location with sub-continental granularity (e.g. North-Europe). One can also define an Affinity group, a group containing a selected set of your services, making sure that the performance is

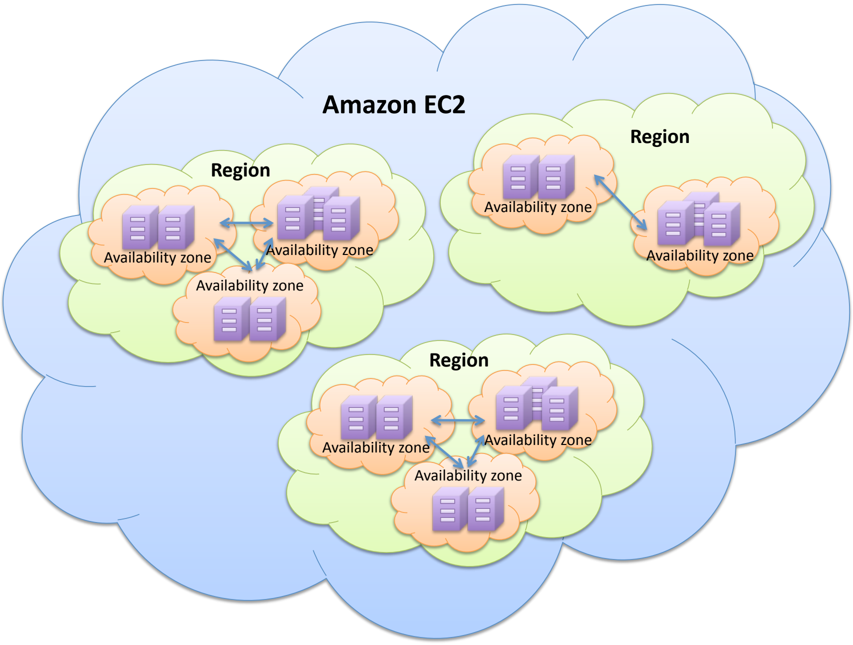

Figure 2.4: Amazon's Availability Zones. Image adapted from [34].

not degraded due to long distances – these Affinity groups can also be geo-located. What happens to the data residing on a disaster-struck data center, is not described in detail, but since legal issues are one of the pros of the geo-location service, stated by Microsoft themselves[33], one can make the assumption that the data will not reside outside the area specified even if a data center is made unavailable. Details regarding location and Affinity group can be found programmatically through the *Service Management REST API*.

Amazon Web Services

Amazon also allow you to specify locations for your services. As can be seen in Figure 2.4 the AWS EC2 offering is divided into *availability zones*. These zones are comprised of *regions* [34]. Each region is designed to operate completely independent of the others to ensure continued service. Of course, if customers choose to keep all their EC2 instances within one single region it will be rendered unavailable in the case the data center suffers failure. AWS S3 (storage) instances are confined to specified regions as well.

Google AppEngine

Google AppEngine does not allow you to specify a location for your services. As earlier mentioned, App Engine provide compute and storage resources, not a virtual machine running in a data center. Even though the specific algorithms and routines for resource pooling are unknown, it would be likely that resource pooling occur since no locality promises are made. If no location data are provided when uploading the program, Google does not know where you would prefer it being resident.

2.4.9 Internal Access Control

When an administrator from the provider access the data on the hosting operating system, this must be thoroughly controlled and logged to prevent any unwanted changes to data and client instances. A provider administrator should not have access to the guest Operating System (OS), the hosted system. Any access to the guest OS should be logged and one should be able to disable password-based login – utilising only token- or key-based authentication. This is also valid for outsiders claiming unauthorised access to a system, either logically (electronically) or physically (by showing up).

Windows Azure

The Microsoft Global Foundation Services (GFS), which manage the cloud infrastructure[10] and platform for Microsoft online services[11], employs a least-privilege and need-to-know policy model for all employees affiliated with the cloud technologies [29]. This would also include Azure.

In a section stating how Microsoft applies controls to physical security the *Securing Microsoft's Cloud Infrastructure* whitepaper states that:

> *Access is restricted by applying a least privilege policy, so that only essential personnel are authorized to manage customers' applications and services.* – GFS

This whitepaper is primarily covering the Microsoft cloud, not specifically the Azure platform – which is hosted on the Microsoft cloud. This would lead us to believe that this description of physical security would also be valid for data hosted in Azure, meaning that essential personnel could have access to the the customer compute instance and data.

[10] The GFS-managed infrastructure include data center facilities, hardware and components supporting services and networks.

[11] The platform supporting the Microsoft online services include compute runtimes such as web server software, SQL servers, identify and directory stores, name services and other consumable functions

Amazon Web Services

According to Amazon [20], they have implemented a two tier security. One security layer for the AWS administrators, accessing the host OS, running directly on the data centers; one layer for the guest OS running on virtual machines. These guest OS are the systems the users get access to.

To access the host OS, the Amazon administrator mush use an individual cryptographically strong Secure Shell (SSH)-key to access the bastion host[12]. Once authorised by the bastion, the administrator can escalate privileges to gain access to a specific host. All accesses are logged and audited on a routinely basis. The privileges to administrate a host is distributed on a "principle of least privilege"-basis, which means that after the administrative tasks on the host are complete, the privileges and access to bastion host are revoked.

The virtual guest OS is completely under the control of the user. The user will have full privileges on the guest host, with root access to the operating system, full administrative privileges over applications, services and user accounts. The Amazon administrator has no access to the guest OS and can not help the user with strictly administrative issues without a hand-over of control. Many operating systems provide the possibility to disable password-authentication, forcing key- or token-based authentication, eg. with a SSH-key pair on a UNIX system. Also, one would be able to log all root access, by forcing use of the `sudo` utility for privilege escalation and enabling shell logging.

For their data centers, Amazon claims to have military grade, multi-tiered physical access control and security. The exact location of the data centers is never disclosed to non-pertinent personnel. Both perimeter and building checkpoints must be cleared before entering the facilities. Amazon claims to have no less than three separate two-factor authorisation checkpoints. All visitors to the facilities are strictly controlled and are at all times followed by authorised staff. Physical access is, for all Amazon employees, issued on a business objective basis, which means that privileges are immediately revoked when the employee no longer has business needs within the facility. All physical access is logged and audited, as with data access.

Google AppEngine

> *(...) Google is able to hire many of the worlds leading security experts to protect our systems and conduct cutting-edge security research. Our data centers are hardened with many of the latest measures in security precautions, including biometric access controls and multi-tiered security perimeters. Furthermore, Google has implemented a multi-layered security process protocol designed to help keep customer data safe. Our*

[12] A bastion host is a computer fully exposed to attacks, but is design to withstand intrusion attempts. This specialised machine is located on the public side of a Demilitarised Zone (DMZ) and is used to access and protect the management plane of the cloud.

> *processes have been independently verified in a successful third-party SAS 70 Type II audit to verify our confidentiality, integrity and availability of customer data.* – Google Apps [35]

As one can see, this quote comes from Google Apps, but since the AppEngine's datastore is powered by mainly the two Google services; Bigtable and Google File System – both used by an array of Google Apps (such as GMail and Docs), this leads us to conclude that the published security statements are also valid for the AppEngine, since the services used are the same and the data is hosted in the same data centers – covering both logical- and physical access. The information presented in this particular section will, to a large extent, be based on the knowledge present about Google Apps data center security from a Google whitepaper – *Comprehensive review of security and vulnerability protections for Google Apps* [36].

Many of the same statements regarding physical data access are valid for Google as the two other discussed providers; the latest in security precautions at the data centers; multi-factor, multi-layered authentication for access to facilities and server containers by selected Google personnel.

With regards to logical security, Google has enabled a series of in-depth security measures to prevent unauthorised access from third parties. Special purpose software, both proprietary and commercial software, is used to protect the system. To prevent security breaches through low level programs, such as interpreters and compilers, these are modified and secured in multiple layers. Not much is guaranteed when it comes to Google employee access to the logical data other than the restricted access to specific employees and the processes defined and validated according to internal auditing. Some of the data center processes are verified in a SAS 70 Type II audit and include both physical- and logical security handling. This is more thoroughly explained in Section 4.3.1.3.

2.4.10 Virtualisation

For multi-tenant cloud services to be provided at infrastructure, platform or application level, the use of virtualisation is inevitable for the provider to remain cost efficient. As stated in Chapter 2, virtualisation, together with grid computing, are the backbone of cloud computing as we normally refer to it. There exists several abstractional variations of virtualisation, but in this section, hardware virtualisation will be addressed. Hardware virtualisation is from now on referred to as virtualisation. The areas of focus within this subject is shown in Figure 2.5.

Windows Azure

Little is known about the hypervisor used for Windows Azure, and information has yet to be reported through any official whitepapers. Some information has been provided by Microsoft through a mail sent to Keith Ward, the Editor of the publication *Virtualization Review*[13] [37]. It is reported that Windows Azure runs

[13] See `http://virtualizationreview.com`

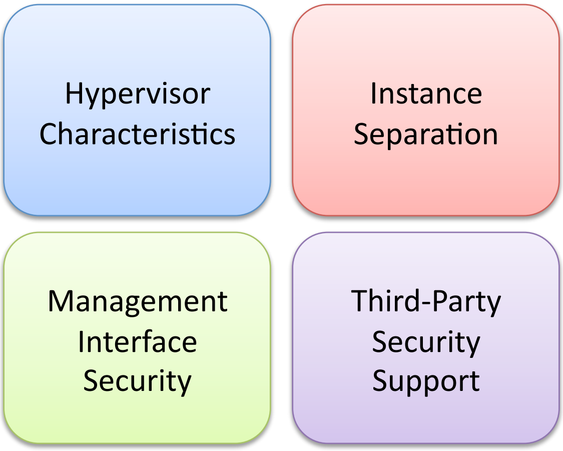
Hypervisor Characteristics
Instance Separation
Management Interface Security
Third-Party Security Support

Figure 2.5: The virtualisation security areas of focus.

on a modified Hyper-V hypervisor, and that this version won't be available commercially due to the advantage it takes from the specific, homogenous data center environment running at Microsoft[38]. The hypervisor is connected to a dedicated fabric controller, ensuring the scalability required for such a cloud environment.

The Virtual Machine (VM) and Hypervisor are transparent to the user, and both virtualisation and OS management is abstracted away from the user, enabling full focus on the deployed application rather than the hardware. *Virtualization Review* [37] also reveals Microsoft's plans to offer Windows Server virtual machine support for Windows Azure, to make it easier for customers to support virtualised infrastructure spread over both cloud instances and on-premise centers.

With regards to compartementalisation and strength of the VM, many of the weaknesses identified for Hyper-V stems from fact that each host has an underlying Windows OS. Virtualisation and OS live side by side, as shown in Figure 2.6 along with a similarily featured AWS instance. If one were to look at the "Critical Security Fix"-release frequency and all the general scrutiny Windows receives, one can assume that critical weaknesses still exist. Windows is a multipurpose OS, and adding the hypervisor role obviously doesn't make it less heavy and more secure. Tom Bittman, VP and analyst with Gardner believes that running all I/O controls and other managers through the parent OS will make the architecture prone to "a single point of failure".[39] Although this may be valid for the commercial hypervisor, it may have been mitigated when customising Hyper-V for Windows Azure. Windows Azure is not like the average, multipurpose OS, and is specially designed to run virtual instances on a grid network – major changes have been done, and strengthening security before launching the cloud environment might very well have been one of them. Microsoft Online Services teams apply the Security Development Lifecycle (SDL) principles which includes *security in development, security by default and security in deployment* [40], and if these are followed, this would

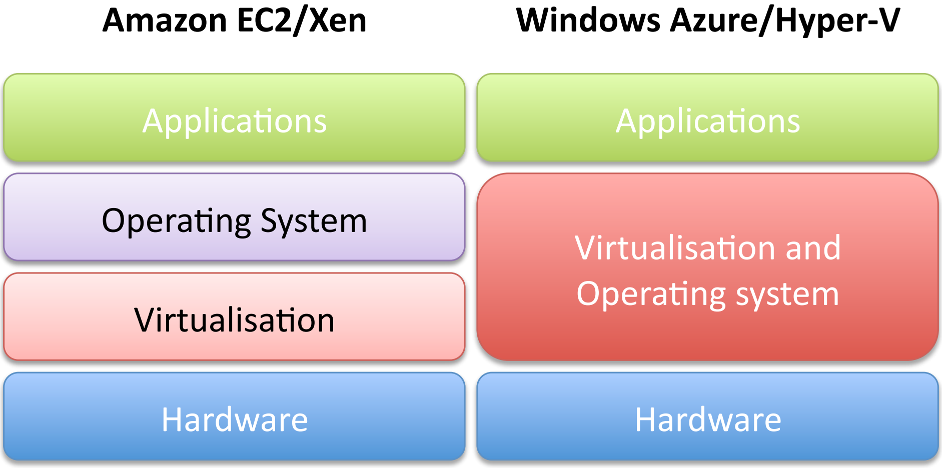

Figure 2.6: The figure shows how the virtualisation and operating system comprise the virtual machine environment.

strengthen the indication of a more secure Azure hypervisor.

There is no reason to believe that Windows Azure include any third-party security software, such as firewalls or virus scanners. If highly regarded security software was included as a part of the solution, the user would not have to put all trust in, and rely on, one provider without specific security domain knowledge.

Web based management tools are protected by X.509 public-key infrastructure and transport security is ensured through SSL. All management API are accessed over REST services, with the same security features as the web based interface [41]. The Azure administrative web interface is protected by a one-factor password authentication over SSL. In September of 2009, Zend[14] announced the *Simple Cloud API* project[15]. This Open Source Software (OSS) project was co-founded by Microsoft, IBM and others. This PHP: Hypertext Preprocessor (PHP) API is intended to be an abstraction layer between the storage services of the cloud providers. To date, both Microsoft's and Amazon's storage services are supported. More projects like this are likely to surface in the coming years. It is conceivable that the use of such APIs will likely enhance the security of current and future applications by assisting customers in their usage of the cloud vendors' storage APIs.

Amazon Web Services

AWS uses a Xen hypervisor. Each instance is called a *virtual private server*. The architecture is depicted in Figure 2.7. On top of physical layer – the Amazon data

[14] See `http://www.zend.com/en/`
[15] See `http://simplecloud.org`

centers – the user can specify security rules for incoming requests.

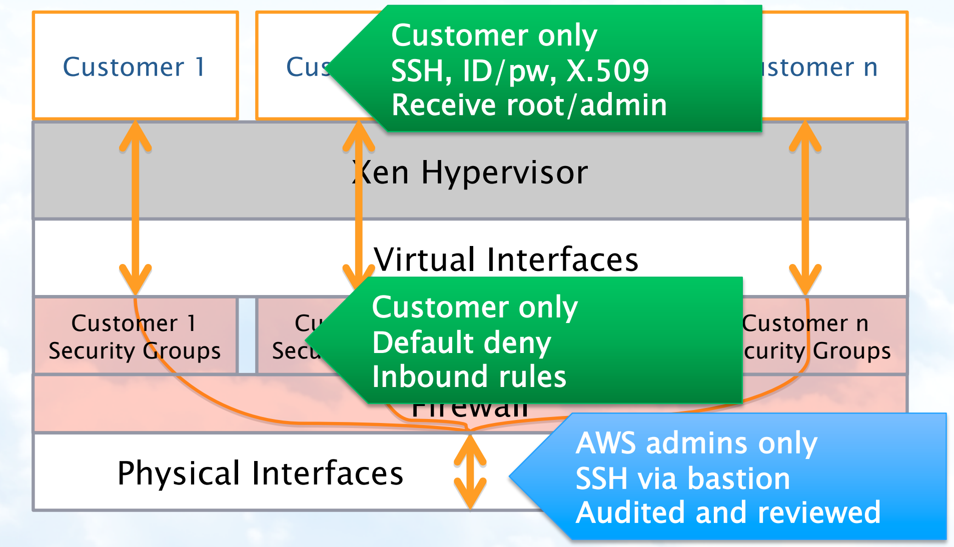

Figure 2.7: AWS architecture shown together with security measures in customer instances, security rule management and physical layer. (enStratus and Amazon, 2009)

A highly customised version of the open source Xen hypervisor (from now on referred to as *the hypervisor*) from Citrix provides the virtual machine hypervisor environment to Amazon Web Services. Amazon is active in the contribution to the project [20]. Historically, Amazon has had their technical focus towards selling products through a large internet store residing on Amazon data centers, and not providing virtual machine technology. Citrix is a company specialising in virtualisation, networking and cloud in general. The Amazon cloud can be augmented with technology from third-party vendors, such as enStratus[16] and different AMI[17] providers, such as TurnKey Linux Virtual Appliance Library[18] and Bitnami[19] – each providing security measures at the relevant level of abstraction.

The Xen hypervisor takes advantage of *paravirtualisation*[20]. Since calls to the underlying hardware requires elevated privileges from the user, it is – with the paravirtualisation scheme – possible to run the guest OS without any elevated access to Central Processing Unit (CPU) routines. The complete virtualisation

[16] enStratus provides an interface for accessing AWS data, a two-step encryption scheme, backup scheduling and encryption and a process for secure AMI bundling.

[17] The basic unit of deployment to the Amazon EC2 instance.

[18] See `http://www.turnkeylinux.org/`

[19] See `http://bitnami.org/`

[20] The hypervisor presents a software interface to access the underlying hardware, similar to, but not identical to presenting direct access to underlying infrastructure and hardware.

of the physical resources, results in a security separation between guest OS and physical hardware [20].

One theoretical weakness with the hypervisor that may plausibly exist is a sandboxing issue. It involves injecting a malicious virtual machine on the same physical location as the target, and then attacking the victim.[2] To identify the physical location of any target VM, one can simply start a rather large set of VM's randomly, due to the low cost of starting new instances compared to the cost of scaling them later. Talbot has estimated a 40 % chance of landing a VM on the same location machine as the target VM. The attacks will mostly consist of stealing encryption keys, key-stroke attacks and other attacks related to reading the disclosed data. No reported successful attacks employing this method are published.

The separation of the virtual instances is handled by the hypervisor. As depicted in Figure 2.7, all traffic must pass through the firewall in the hypervisor layer. The rules defined in this firewall defines who can access the instance, limiting access to only authenticated users.

Physical disks are also virtualised to provide virtual disk space for the user. Since many users share the same hardware, disks are completely wiped prior to other users gaining access. This ensures that one customer's data will not be accessible to other customers, even though the same hardware is being "reused". The physical RAM is partitioned in a similar manner as the hypervisor separation of instances, described in the previous paragraph.

Calls to the administrative interfaces of the instance must be signed by a X.509 certificate or by the proprietary *Amazon Secret Access Key*[21]. Using Secure Socket Layer (SSL) for API calls are not required, but this option is easily available and recommended by Amazon. Also, a number of security related network features are included by default [20]. The administrative web interface is protected by a one-factor password authentication over SSL.

Google AppEngine

> *I think it will be very sad if we need to use virtualization. (...) It is hard to claim we will never use it, but we don't really use it today.* - Luiz André Barroso, Google

Previously in this section, virtualisation has been defined as a hypervisor providing virtual machines, or virtual instances of some sort. According to this definition, and the statements of Google engineer Luiz André Barroso, Google does not offer (hardware) virtualisation.

Although this not being strictly relevant for this section, application separation and application management security measures will be discussed here, due to the comparative interest for these areas and the similiar Amazon Web Services and Windows Azure virtual machine separation and management API security. The GAE management console is exclusively provided over SSL and authentication is

[21] When signing up for an AWS account, a Secret Access Key is issued to the customer.

done with a Google federated identity, a one-factor password authentication also used for other Google products. Authentication is done over forced SSL.

Since all user applications are run in a non-specific environment, Google does not allow the user to access the file system for file storage. All persisted data must go through a storage API. Since both processing and data is made transparent through the abstraction API, not much information can be derived from wether or not the data is separated, but since all storage and processing are done through a strictly, programmatically accessible API, the possibilities of vulnerability exploitation are minor. One of the most talked about might-be issues, are problems identified in the Python interpreter[42]. These issues might be translated to the interpreter running on the App Engine. Guido van Rossum states in an interview that these issues are being taken care of, by both sandboxing each user instance, selecting which core modules are to be supported in the Python environment – based on some security criteria – and hardening the interpreter itself. In addition, more in-depth security measures are implemented, but more detailed information on these measures are disclosed.

According to Google, by April 6th 2010, their AppEngine has had over 250 million page views and the community consists of over 250 thousand developers[43]. Even with this usage, few issues have been identified, let alone published for the world to see.

2.5 Summary

> *The key takeaway for security architecture is that the lower down the stack the cloud service provider stops, the more security capabilities and management consumers are responsible for implementing and managing themselves.* – Cloud Security Alliance [3]

The security implications of the services built upon the different platforms can in this non-quantifiable way be measured by how much of the infrastructure is available from SaaS through PaaS to IaaS. Figure 2.2 showed that the cloud vendors do not necessarily stay with a single model. While Microsoft has kept to the PaaS model, both Google and Amazon offer services within two models. These differences due to feature complexity are not only apparent between the three service models but also between vendors within each model.

Are there any major differences between the vendors?

- **Service Level Agreements and Compliance Features:** The only differences to speak of are the tiny differences in the amount of discount the customer is entitled to.
- **Authentication Services:** Microsoft's Access Control Service, part of its AppFabric offering, enables very easy use of federated identity services across companies in a secure manner. Neither Amazon nor Google offer the same functionality to the same extent.

- **Audit and Certifications:** All vendors are both Type I and Type II SAS70 certified, but only Microsoft is ISO 27001 certified to date.
- **Cloud Tech Support:** No real differences to speak of. All vendors provide 1:1 support around the clock support at differing price levels.
- **Incident Response and Logging:** Nor here are there large differences. Google offers somewhat less accessible information since it is a SaaS vendor.
- **On-demand Self-service:** Both Microsoft and Azure offer APIs for programmatic scaling and easy control over the customers instances. Google merely allows an on/off switch and handles issues related to scaling.
- **Broad Network Access:** There seems to be a consensus across the surveyed vendors to provide RESTful APIs, thus enabling the consuming of services from any platform in an open and standardised way.
- **Resource Pooling:** Both Microsoft and Amazon have come to the conclusion that the customer should be allowed to limit the geographic dispersion of their applications and data. To date, Google does not offer this control.
- **Internal Access Control:** All surveyed vendors claim to take access control, both physical and virtual, very seriously. They have implemented military grade perimeter security and access auditing.
- **Virtualisation:** Google AppEngine differ from the other two platforms by not providing traditional hardware virtualisation, and therefore not giving any notion of a virtual machine or the possibility to administrate this for the user. Amazon Web Services and Windows Azure are more similar in nature, albeit Azure relies on a completely proprietary platform compared to Amazon's Xen hypervisor and the possibility for including third party software. Due to the lack of instance control, AppEngine is not a regarded as a good solution for the dSafe application.

What do the differences just summarised say about these vendors? There does not seem to be any large differences. After trying to locate information and specifications for several months, it has become our experience that the surveyed vendors all try to be open and transparent and still they maintain many business secrets. This gets in the way of potential customers's security needs. Companies considering going into the cloud will need to continue pushing for transparency if they want to be guaranteed the security of having their own private data center.

In this comparative study three PaaS vendors were surveyed: Microsoft, Amazon and Google. As might be expected from vendors that compete in the same PaaS segment, their feature sets seem to be *converging*. Each vendor tries to become par with the other vendors while also creating that new feature to give them that elusive business edge. There are three good examples of this: First, Microsoft's new platform has taken a lot of cues from AWS by facilitating programmatic scaling of applications as well as allowing the setting of affinity groups. Second, Google has very recently decided to create a storage solution that specifically targets Amazon and Microsoft's current storage services. And third, all the surveyed vendors are at a consensus of supporting RESTful web services as communication interfaces.

Chapter 3

Case Overview

This chapter contains a description of the start-up business dSafe. We analyse their business plan as well as their design documents to find implications for system security. dSafe always intended cloud deployment of their system, specifically on Microsoft's Azure platform. We therefore intend to conclude whether dSafe's application can be adequately secured on this platform. From this security perspective, both legal and practical angles are taken into consideration.

This chapter contains the following sections:

- **The Company:** Details regarding the company such as leadership structure and the reasons for founding it.
- **Information Flow and Storage:** This section contains diagrams, descriptions and explanations of dSafe's planned architecture and information flow. The section also contains information regarding the nature of the information they intend to store in their databases.
- **Technical Challenges:** This section contains a quick analysis of what technical challenges dSafe faces in its design and implementation. The challenges, technical and otherwise, related to cloud technology are not covered until Chapter 4.

3.1 The Company

dSafe is a company developing a secure digital safety deposit box for storing receipts generated by credit card use, contracts or any other information that is suitable for storage within banking information systems. The company originates from the Entrepreneurship Center at the Norwegian University of Science and Technology and was started by Daro Navaratnam, the current Business Developer at dSafe.

The company was established in 2009, with the goal of creating an online, secure and centralised storage for retail receipts – simplifying the process of warranty claims and other purchase annulations or refunds. To achieve this, cooperation and integration with retail stores are essential. Partnership with large department chains, focusing on expensive goods with extended warranty time, will be the primary focus in early stages of the business plan execution. By the end 2010, dSafe aims to be fully integrated with 1000 retail stores before launching the product officially. dSafe has already acquired partners within the sports industry, and intends to bring in electronic and grocery chains. By the end of 2011, the aim is to become fully integrated with current web-based banking solutions and to provide the receipt storage through existing internet banking interfaces.

dSafe is now comprised of 5 employees, 3 full-time and 2 part-time, whereof 4 are programmers. All employees have technical academic backgrounds from the Norwegian University of Science and Technology.

3.1.1 Business Plan

Navaratnam has developed a plan for dSafe's progress. Figure 3.1 shows an overview of the planned implementation of the software assets realising the business concept. The Gantt diagram has not been translated to avoid any informational loss during the transformation[1].

Problem

Most consumers receive a paper receipt when purchasing goods, but the organisation and storage of all receipts are comprehensive tasks. Therefore paper receipts are easily lost. This can be a problem, since the onus of proof lies on the consumer side of the purchase when returning or claiming a refund of goods. dSafe claims that over 52% of the consumers have experienced lack of service and difficulties associated with making a claim when failing to present a valid receipt. A receipt is needed in the following cases:

- Return and reclaim of goods (at most 5 years after purchase).
- Proof of purchase in insurance claims.
- Refund of travel expenses.
- The Norwegian Accounting Act states that receipts posted in the accounts must be stored for at least 10 years.

Solution

The service *dKvittering* (dReceipt) will make sure that all consumer receipts are automatically sent from the store to the dSafe storage, where the consumer will have access to all receipts online through the dSafe consumer portal or via an internet banking service. The consumer will be identified through the bank card used for the transaction. dKvittering can also provide services such as:

[1] Readers are referred to Google Translate.

- Warnings to consumers that has purchased goods later identified as hazardous or goods being recalled by the producer.
- An automatic overview of personal expenses and economy.
- Electronic supplements and vouchers for accounting.
- Easier and more accurate keeping of travel expenses.

3.1.2 Market

Web based banking solutions and electronic payments are popular among Norwegian consumers as internet has grown to become accessible by the majority of the population. During the first phase dSafe will implement the receipt storage service for expensive luxury goods, such as sporting goods, electronics and furniture. In Norway, these business segments are represented through large chains which makes it easier to implement a solution towards an array of stores. The international market is harder to predict, since a Danish company released a similar solution, called *ekvittering.dk* in April of 2009. Also, a start-up company in Sweden is working towards offering digital receipts as a service. The market potential in Norway alone is estimated to approximately 1 million users, based on the amount of people already using electronic payment services.

3.1.3 Business Model

Since the burden of proof lies on the consumer when purchasing a commodity, the consumers will have the greatest need for such a service. The market analysis shows that most people are willing to pay approximately 15 to 20 NOK per month to avoid having to store and organise paper receipts. dSafe will also profit on selling market statistics. To enter this market quickly, dSafe can cooperate with other businesses, e.g. Norstat, already present in this segment.

3.2 Information Flow and Storage

In order to locate any weaknesses one must identify the information which is protected by the software. This section will describe the information flow of the dSafe information system and what information is persisted in the system. Local laws, regulations and norms will also be taken under consideration when auditing the information.

3.2.1 The data flow

dSafe's current plans for their architecture is depicted in Figure 3.2. It contains a number of components:

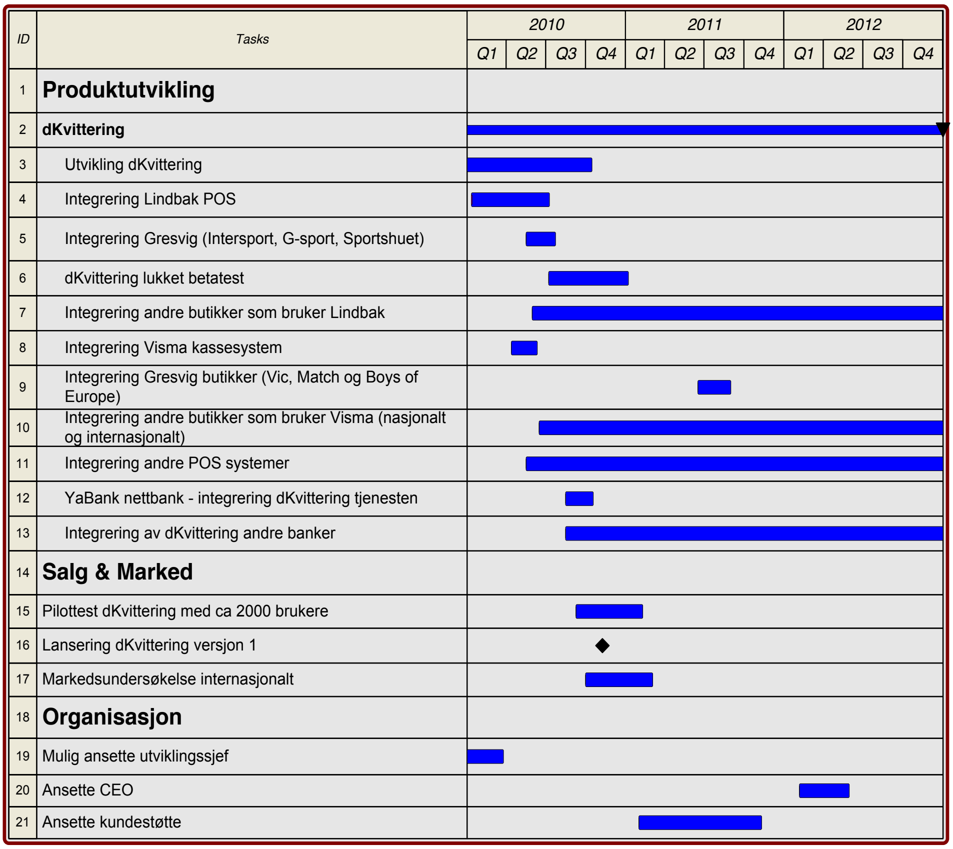

ID	Tasks
1	**Produktutvikling**
2	**dKvittering**
3	Utvikling dKvittering
4	Integrering Lindbak POS
5	Integrering Gresvig (Intersport, G-sport, Sportshuet)
6	dKvittering lukket betatest
7	Integrering andre butikker som bruker Lindbak
8	Integrering Visma kassesystem
9	Integrering Gresvig butikker (Vic, Match og Boys of Europe)
10	Integrering andre butikker som bruker Visma (nasjonalt og internasjonalt)
11	Integrering andre POS systemer
12	YaBank nettbank - integrering dKvittering tjenesten
13	Integrering av dKvittering andre banker
14	**Salg & Marked**
15	Pilottest dKvittering med ca 2000 brukere
16	Lansering dKvittering versjon 1
17	Markedsundersøkelse internasjonalt
18	**Organisasjon**
19	Mulig ansette utviklingssjef
20	Ansette CEO
21	Ansette kundestøtte

Figure 3.1: Gantt-diagram over dSafe activities

- Two web services: One for stores uploading receipts and one for users to access their data.
- One processing application: This performs parsing and storage.
- Two databases: One for customer data and a separate one for receipt data.
- Two raw storage devices: One for queued receipts and one where said receipts are stored post processing.

The main data flow of the system is described below. Only the correct operational pattern is included, although Figure 3.2 contains components that pertain to what happens when certain operations fail.

1. A partner retail store consults a web service to find out whether or not the credit card being used is associated with one of dSafe's users.
2. If this is found to be the case, the store sends receipt data to another web service over a secure, authenticated connection.

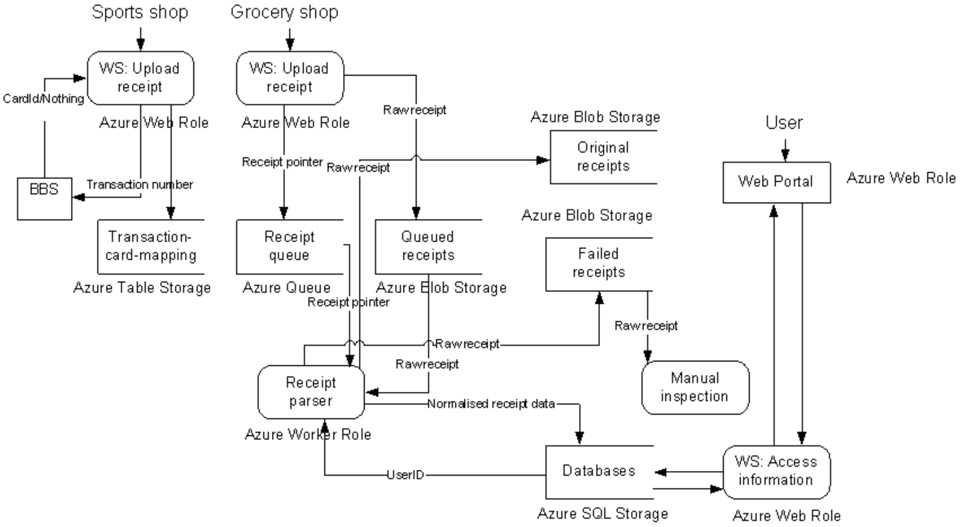

Figure 3.2: Top level Data Flow Diagram

3. The web service generates a hash of the file, does a check to see if the hash is unique in a dictionary. If the hash is unique, the incoming data is added to a blob storage container named `Queued Reciepts`, and a pointer to the file is added to a queue named `Receipt Queue`.
4. An application polls from the `Receipt Queue`. It fetches the corresponding data from the `Receipt Queue Storage` and sends it to a corresponding parser module.
5. The parser module reads data from the raw receipt data and returns a receipt object (Data Transfer Object (DTO)) with all the data that is going to be stored in the database.
6. The module sends the DTO to an Object-Relational Mapping (ORM) that in turn inserts the receipt data in the receipt database.
7. The worker role copies the original receipt data from `Receipt Queue Storage` to a separate blob storage named `Original Receipt Storage`.
8. Finally the worker role copies the original receipt data from `Rececipt Queue Storage` to a separate Azure Blob Storage container named `Original Receipt Storage`.

3.2.2 Persisted storage

From the definitions in Section 2 of the Norwegian Personal Data Act[44]:

> *1) personal data: any information and assessments that may be linked to a natural person*

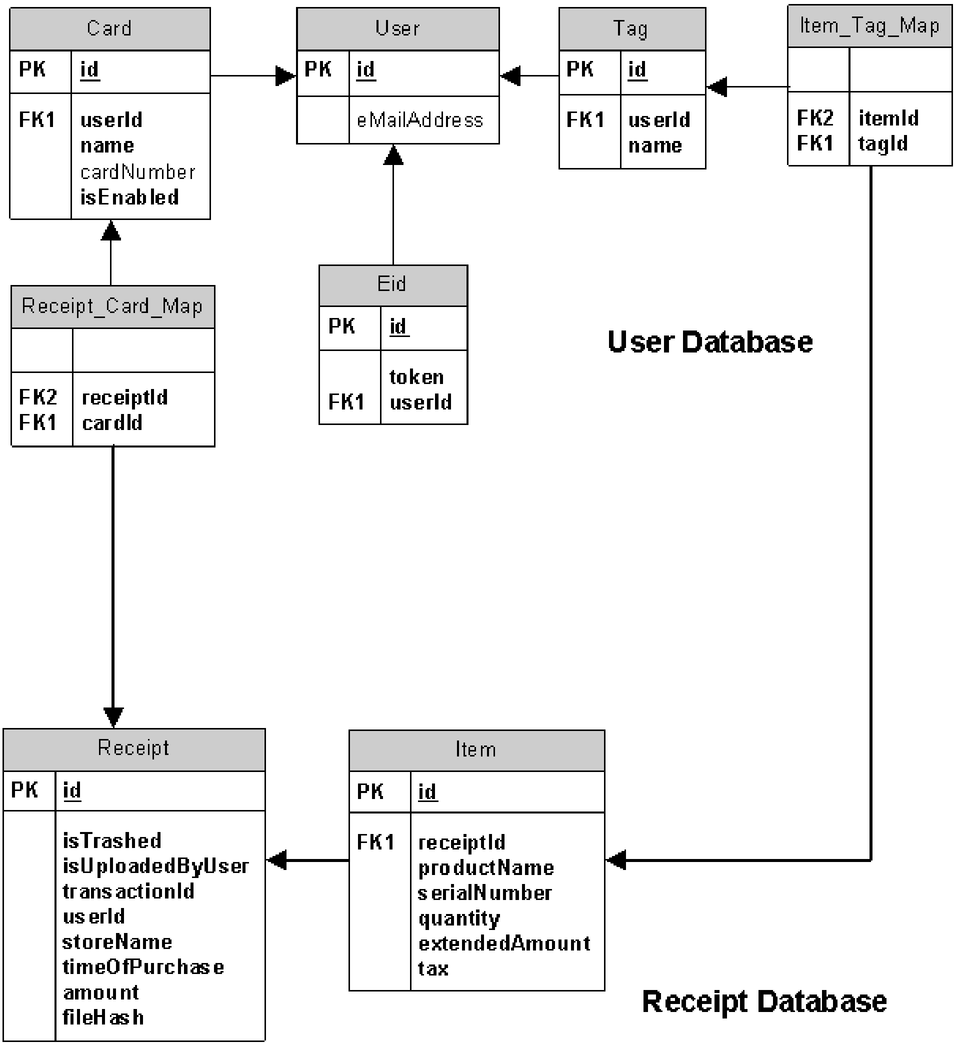

Figure 3.3: ER diagram of dSafe's databases

> *8) sensitive personal data: information relating to a) racial or ethnic origin, or political opinions, philosophical or religious beliefs,*
> *b) the fact that a person has been suspected of, charged with, indicted for or convicted of a criminal act,*
> *c) health,*
> *d) sex life,*
> *e) trade-union membership.*

dSafe intends to store:

1. Receipts, which includes data on the specifics of each purchase.
2. Categorisation of receipt contents and customisable reports.
3. Contracts and their digital signatures (planned feature).

The detailed receipts of individuals tells a tale of behavioral patterns, revealing both movement and shopping patterns. This could be very useful information to both advertisers and criminals. E.g. identity thefts are mostly financially motivated and the receipt data provide an excellent way to locate wealthy victims [45].

All of the information dSafe intends to store is clearly personal information as it pertains to persons (satisfying Section 2 definition point 1 in said act), and we argue that both receipts and contracts may contain information that can be regarded as sensitive. For instance, receipts could come from sex shops (point 8d) or pharmacies (point 8c).

3.3 Technical Challenges

The legal challenges previously described imply several technical challenges. According to Figure 3.2 there are two external entry points. In this section, the web service that is to be used by the stores is referred to as E1 (*External 1*) and the one that is to be used by end-users as E2 (*External 2*).

3.3.1 Confidentiality

The confidentiality of the data passed between dSafe's services and client applications must be protected to prevent information from falling into malicious hands. The integrity of the receipt data is paramount, as digital receipts are not only intended as a convenience but also as proofs-of-purchase for warranty and insurance claims. Using SSL to encrypt messages is a long standing best practice in the security community and is described in the security pattern *Secure Pipe*[46, chap. 9] and as such, it is a good solution to this problem. dSafe intends to implement this pattern in both external entry points.

3.3.2 Identification and Authentication

For both external entry points, the interacting entity, be it end-users or stores, must be properly identified and authenticated to prevent untoward information disclosure. dSafe intends to employ username and password credentials over an SSL secured connection to achieve these properties. If, in this rudimentary authentication scheme, it could be guaranteed that these credentials are never transferred unsecured, this identification scheme could suffice. Thus, dSafe would be able to maintain which store uploads which receipts. However, since this upload operation is very important to all involved it might be more appropriate to use some more elaborate and more reliable authentication service.

3.3.3 Non-repudiation

In E1, the stores upload receipt data. The web service will need to enforce non-repudiation, so that stores are unable to deny that a receipt upload has taken place. Otherwise, the digital receipt can never become equivalent to the paper version. This will ensure that the users can trust dSafe's service enough to discard paper receipts, thus achieving the goal of furthering the "paperless society".

On the other hand: To protect stores the service will need to provide *at most one*-atomicity so that receipts are never duplicated in the database. As can be observed in Figure 3.2, the service will hash received data and check for duplicates. The hash algorithm will ensure that identical receipts are ignored.

Chapter 4

Case Study

In this chapter we will discuss how dSafe's application can be deployed securely in a cloud. To be able to recommend an approach for cloud deployment, a defined process is followed in order to identify the most valid risk scenarios for the deployed assets. The process was defined in Section 1.2.

This chapter contains the following sections:

- **Analysing the Cloud Provider Risk:** In which we analyse the architectural designs of dSafe that were presented in Chapter 3. We identify the different assets of dSafe's system and evaluate the risks related to putting each of them into the care of a cloud provider.
- **Deployment Model Acceptance:** This section covers what deployment model dSafe should use to meet both business-related functional requirements and to facilitate the mitigation of risks previously identified.
- **The Twelve Domains of Critical Focus:** The Cloud Security Alliance describes these domains in detail. In this section we put them into dSafe's context and make concrete recommendations as to how dSafe should handle the many decisions prompted by the domains.
- **Risk Mitigation:** In this section mitigations are extracted from the discussion involving the twelve domains. The section details the reasoning for each domain and summarise the capabilities of respectively Microsoft, Amazon and Google.

4.1 Analysing the Risk

This section is dedicated to identifying the assets of the dSafe application, evaluating the associated risks and finding mitigation strategies. We utilise the Cloud Security Guidance framework provided by the CSA which is meant for companies considering the possibilities of moving their applications into the cloud. It provides a risk-based approach, but the process is not in any way intended to fully replace a complete risk assessment methodology, nor is it meant to map all security requirements – it is used to evaluate assets' tolerance for being moved into the cloud [3]. In other words, we are evaluating *cloud provider risk* – not any risks present due to programmatic errors or malpractices, or even design level security flaws.

The Oxford Dictionary of English defines the word *asset* thusly: "an item of property owned by a person or company, regarded as having value and available to meet debts, commitments, or legacies". In our context assets are typically databases and web services; items that need protecting – items that are critical for the application and the company. For each of these assets we examine the relevant scenarios that might occur in production, ascertaining the implied consequences. By taking another look at the available cloud provider technologies we try to discover which providers have implemented technologies for mitigating the charted risks. Thus, conclusions on which provider is best suited for dSafe can be reached.

4.1.1 Identifying the Asset for Cloud Deployment

To assess the deployment of the dSafe application to the cloud, one must first identify the assets. The security guidance suggests splitting the assets into two categories:

1. Data
2. Applications/Functions/Processes

These assets should be evaluated separately, identifying which parts are moved to the cloud. Scope creep should also be accounted for, since data and transaction volumes often are higher than expected. One must therefore, in addition to identifying what data and functions should reside on the cloud, take into account *future* use of the data and functions. These assets are categorised in the *scope creep* domain.

Data assets

Databases and other storage of vital data is covered in this section. Table 4.1 lists the identified data assets. Each logical separation of an amount of related data is called an asset. For structuring and ease of localisation, each asset is categorised in a domain, in this case related to the type of storage. Each asset is assigned an ID for later references.

Process assets

Transactions, processes and available functionality of the application falls under the *process asset* category. E.g. worker roles and web services are both process assets. These assets are different in nature, or work on data sets of dissimilar nature. The process assets are listed in Table 4.2 and like the data assets, each process asset

Table 4.1: Table shows all identified data assets with descriptions.

ID	Domain	Asset	Description
DA1	Database	End-user data	Contains all information necessary to achieve billing for dSafe's rendered services.
DA2	Database	Receipt breakdown data	Contains the normalised data gathered from uploaded receipts.
DA3	Unstructured	Pre-processing storage	Raw receipt data uploaded by dSafe's partners.
DA4	Unstructured	Post-processing storage	Backup storage of the unprocessed receipts.
DA5	Unstructured	Failed receipt storage	Storage of receipts which syntax is unrecognisable
DA6	Queue	Receipt pointer queue	Contains pointers to receipts in DA3.
DA7	Scope creep	Contract storage	Contains contracts and signatures

is assigned a unique ID for easier referencing. Each asset is also placed under a domain – related to the nature of the work being done or service offered.

Table 4.2: Table shows all identified process assets with descriptions.

ID	Domain	Asset	Description
PA1	Web service	Receipt upload service	Through which receipts are uploaded (to DA3) for processing (PA3) and storage (DA4).
PA2	Web service	Data retrieval service	Through which users' data can be retrieved (from DA2) for presentation.
PA3	Batch process	Receipt parser	Receipts are parsed, processed and stored in database DA1 for later retrieval.
PA4	Web service	User verification service	Through which stores can find out whether a customer is a dSafe user.
PA5	Scope creep	Contract signer	Uses secure signature to sign legally binding documents.

4.1.2 Evaluating the Asset

To prioritize and evaluate which assets have the greatest risk a set of questions are to be answered for each of them. The risk is determined by defining a probability and consequence to each of the scenarios for each asset. The quantifiers for these factors are low (L), medium (M) and high (H). Figure 4.1 shows the multiplication table for determining the risk. Table 4.3 shows an overview of all data assets and the risk of each scenario applied to the asset. Table 4.4 shows the scenarios applied to all process assets.

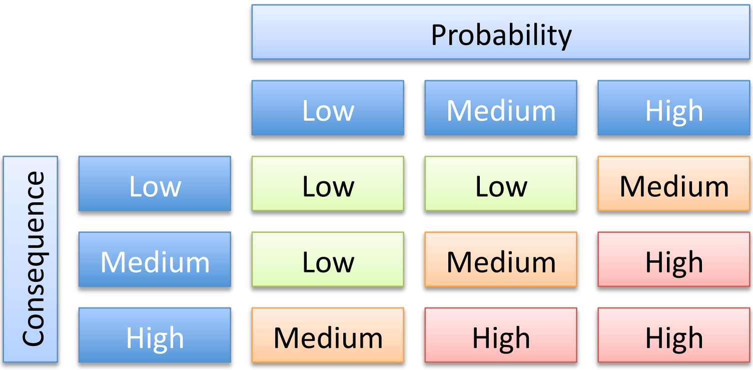

		Probability		
		Low	Medium	High
Consequence	Low	Low	Low	Medium
	Medium	Low	Medium	High
	High	Medium	High	High

Figure 4.1: A risk multiplication matrix that shows risk – annotated by low, medium, high, in the middle of the matrix – being determined by probability and consequence.

The questions to be answered for each asset make up scenarios for each of them:

SC1. How would we be harmed if the asset became widely public and widely distributed?

For a data asset, this would mean that anyone can access the data store and read the data. For a process asset, this means that anyone may use the functionality provided by the process, e.g. calling a web service without needing to bypass any authentication. Everyone, affiliated with dSafe or not, would be able to obtain this information and use any functionality without making an extraordinary effort.

SC2. How would we be harmed if an employee of our cloud provider accessed the asset?

Accessing the asset means reading data from a data asset or using functionality provided by a process asset – the same as the previous question, but restricted to an employee of the cloud provider. An employee of the cloud provider is different from a non-affiliated person due to the fact that the employee might have access to some elevated cloud instance privileges.

SC3. How would we be harmed if the process or function were manipulated by an outsider?

Even though this scenario implies process, the scenario is also valid for the data assets. Manipulated means that the data is changed to another state. For processes, manipulated means that the functionality is changed – the process behaves in another way than normally expected.

SC4. How would we be harmed if the process or function failed to provide expected results?

For a process asset this would mean failing to show correct behaviour, exposing incorrect interfaces, manipulating output or input, retrieving incorrect data etc. This scenario is the equivalent to SC5 for data assets.

SC5. How would we be harmed if the information/data were unexpectedly changed?

For a data asset this would mean failing to keep the integrity of the data over a period of time. This could lead to the later retrieval of incorrect data or the loss of data. This scenario is the equivalent to SC4 for process assets.

SC6. How would we be harmed if the asset were unavailable for a period of time?

This scenario implies that no operation could be done on a data asset. This includes saves, retrievals, data operations or administrative operations. For a process asset this would mean that the process could not be initiated, and no functions would be operational.

For each of the assets identified we now reason about each scenario, determining their probability and consequence. This is combined to become the resulting $[L{\times}H = M]$. In this example the probability is low, consequence high and resulting risk medium.

DA1: End-user information

This database contains the tables described in the top part of Figure 3.3. This database will hold data related to the users of dSafe's services. dSafe will not store usernames and passwords, instead choosing to utilise some sort of unified electronic identity framework. As with the other databases in the system dSafe's design hides their locations and protects them with access control systems.

SC1. dSafe's choice of using an external electronic identity provider greatly mitigates the consequences of exposure. Unless the external database is also exposed the only connection an outsider might find to the user is an email address. In addition, the database contains users' categorisations of rececipts and items contained within; non-critical data. The actual receipts and their broken down data are not contained within this database so the consequence is low. $[L{\times}L = L]$

SC2. Unless the occurrence is publicly broadcasted there will be no loss of trust between the users and dSafe. However, dSafe would have to consider whether this breach of trust and/or contract on the part of the cloud provider is reason enough for sanctions such as discontinuing the contract. $[L{\times}L = L]$

SC3. If the database content were to be altered the worst case is that users might get access to other users' data. Users' categorisation schemes might also be broken. In any case, no receipt data could be lost or manipulated as a direct consequence. [$L \times M = L$]

SC5. Equivalent to SC3.

SC6. No loss of data would occur and stores would still be able to upload new data. Users would not be able to access their data during the time period which would understandably cause irritation of varying degrees. [$L \times L = L$]

DA2: Receipt Breakdown Information

This database will contain the information parsed from the uploaded receipts. These can be very detailed and actually detail exactly which merchandise the end-user has purchased down to the brand of diaper for their kids. It is important to note that this information is close to useless unless database DA1 also has been compromised due to the clear separation of data in dSafe's design. If DA1 were to be compromised, the public would have access to a complete run-down on every purchase end-users complete with their credit cards. If criminals were to use this information to target the end-user as a burglary target this could cause dSafe to be implicated as criminally negligent.

SC1. The primary consequence of this happening will be loss of trust between dSafe and their users since the information is useless to both nosy neighbours and criminals due to the proper data separation described in the last paragraph. The probability of this occurrence is equally low due to the fact that the database's location will not be publicly exposed. [$L \times L = L$]

SC2. Equivalent to DA1 SC2.

SC3. This occurrence could have far-reaching consequences. Not only will it severely undermine the trust between the users and dSafe, but unless the work of the manipulator is discovered and the changes reverted either party (users and stores) could lose money. Due to the storage of original data, a change of this data will not be critical. [$L \times M = L$]

SC5. Equivalent to SC3.

SC6. If this database could not be reached we effectively have a denial of service-situation on our hands. [$L \times H = M$]

DA3: Unprocessed Receipts

This storage will contain raw receipts in differing formats until such time as the process PA3 gets around to retrieving, parsing and moving them to their final resting place, namely DA4.

SC1. It is unlikely that the end-users will see any distinction between having this temporary storage exposed in contrast with the normalised broken down data from data asset DA2 and they would be right. This information is equally useless however, unless both DA4 and DA1 are also exposed. Thus, we come back to end-users not trusting dSafe to keep their data safe. $[L{\times}L = L]$

SC2. Equivalent to SC2 for asset DA1.

SC3. This storage is intended to be kept in synchronisation with DA5 continuously. This means that it is quite simple to detect that raw receipts are deleted. However, there is no way to come back from this. At this point in the data flow this is the only place the receipts' data is stored. $[L{\times}H = M]$

SC5. Equivalent to SC3.

SC6. This would effectively prevent new receipt queuings, but there wouldn't be any loss of existing data due to the downtime. $[L{\times}H = M]$

DA4: Original Receipt Storage

This unstructured storage is intended to contain all receipts as they were transmitted from stores after they have been processed. This could potentially contain more information than dSafe mines from it, for example a partial bank account number.

SC1. The fact that this storage might contain more information than dSafe mines from it would result in higher consequences than for the receipt database (DA2). The probability is still low due to dSafe's design. $[L{\times}M = L]$

SC2. Equivalent to DA1 SC2.

SC3. If this data were to be altered this can be detected by comparing with the hashed value stored in the receipt database. There is no way to recover deleted or altered files though. The consequence is still low, since the parsed data is still contained in DA2. $[L{\times}L = L]$

SC5. This data is not meant to be consulted very often. It is merely a backup database, so it is plausible that it would not be detected unless DA2 suffered a cataclysmic failure. $[L{\times}M = L]$

SC6. Raw receipts could then not be moved here from pre-processing storage, thus disabling process PA3 from fulfilling its mandate of processing receipts. This would halt the system that is concerned with incoming data. The system is intended to recieve and process alot of data, so depending on how long this asset is unavailable the consequences could become dire. $[L{\times}H = M]$

DA5: Failed Receipt Storage

This storage is for receipts that fail to meet process PA3's format expectations. These are stored for manual inspection.

SC1. Equivalent to DA4 SC1.

SC2. Equivalent to DA4 SC2.

SC3. Receipts end up here if they do not fit one of the pre-defined schemes recognisable by process PA3. If outsiders manipulate receipts after they have already been rejected we suggest that they have not achieved anything. However, if the outsider continously deletes receipts before the failed receipts are discovered these could be lost. [$L \times H = M$]

SC5. Equivivalent to SC3.

SC6. This would prevent process PA3 in storing receipts that did not match the recognisable schemes. This would entail that the receipts pile up in queue DA6 but are not lost. [$L \times L = L$]

DA6: Receipt Pointer Queue

This queue contains pointers into storage DA4.

SC1. This data consisting only of relative pointers to files is only useful to process PA3. [$L \times L = L$]

SC2. As with any other asset dSafe would have to decide whether to take action against the provider. [$L \times L = L$]

SC3. This could prevent process PA3 from processing receipts so the consequence is dire. [$L \times H = M$]

SC5. Equivalent to SC3.

SC6. Equivalent to DA3 SC6.

DA7: Contract Storage

This storage unit is part of the planned scope creep.

SC1. These contracts may contain information declared as sensitive information as defined by Norwegian law. The contracts could also contain business information vital to the survival of companies. This scenario could then have dire consequences for both private persons and businesses. dSafe would be liable if this were to happen. As with other data assets, the unintended disclosure of this content will cause users to distrust dSafe. The probability of exposure must be said to be equal to that of this scenario with the other databases. [$L \times H = M$]

SC2. Equivalent to DA1 SC2.

SC3. If legally binding contracts were to be altered or deleted, this can have dire consequences for the parties involved in the contract. One can only speculate on the specifics of this occurrence since we do not know what types of contracts may be stored. [$L \times H = M$]

SC5. For instance, signatures may be removed or corrupted. This would make the data untrustworthy since dSafe needs to guarantee non-repudiation. [$L \times H = M$]

SC6. No data would be lost, but service PA5 would be unable to provide neither the signing service nor the contract upload service. Contrary to storage DA3, DA4 and DA5, this storage will see alot less data and action. [$L \times M = L$]

PA1: Receipt Upload Service

This asset is the entry point (previously described as E1 in Section 3.3) for all businesses and partnering vendors. The service will move incoming receipts into a queue for further processing.

SC1. The asset is protected from unauthorised access through basic authentication over SSL (further described in Section 3.3.2). If this authentication is broken, a peer may inject counterfeit receipts, opening for possibilities to take advantage of this proof of purchase – even though this item is not actually purchased. A breach of security of this web service may also open for more conventional attacks, denying other users the service or corrupting application-internal data. As this web service, in many ways, is the entry point to the application – and must be exposed on the Internet in order to serve external stores and shops – a fair chance exists that this may be the subject of attack from parties with malicious intent. Such an attack would be disastrous for the integrity of the application and business idea. [$M{\times}H = H$]

SC2. If an employee of the cloud provider bypasses the authentication of the web service, this may entail two scenarios. One; the employee takes advantage of the service security breach and injects counterfeit receipts into the system to take advantage of the trust between the store and dSafe application for the benefit of the employee, or a third party. Two; the breach gets publicly known and the trust between stores and dSafe gets critically weakened due to the fact that injecting false data is now possible. [$L{\times}M = L$]

SC3. In addition to the similarities this scenario shares with scenario SC1, a manipulation of this service could lead to corrupting or mining of other users incoming receipt data. If an outsider were to manipulate an incoming message, he could also change the data to be stored to the *original receipt storage* and other, deeper, security measures. The only way to prove the integrity of the data, if questioned, is to compare it with the stores' transmission logs. [$L{\times}H = M$]

SC4. The expected result of this web service is that it should put the raw receipt in an unstructured storage container and a pointer to this receipt into the receipt queue. If the web service fails to store the pointer, the receipt will not be processed and will only be discovered by some sweep of the storage container. If the file itself is not stored, a later process will discover this as a read is performed on the queue. In the case that neither file nor pointer is forwarded to the correct service, the receipt will be lost. If the web service fails to extract the receipt from the incoming message, or does not receive a receipt properly this is a repudiation issue, that must be solved together with the sender – the store. [$L{\times}H = M$]

SC6. If the web service is unavailable, no incoming receipts are recorded. The burden of ensuring full logging of all purchases is therefore also in the hands of the sender. Serious programmatic errors, cloud service downtime and other network related availability issues may lead to this scenario – at least for a short period of time. [$M{\times}H = H$]

PA2: Data Retrieval Service

The web service from which web portals, banking solutions or other third party solutions may fetch data for presentation. The consumer must authenticate against this interface with some kind of electronic ID. The service produces receipt history, -details and user data from the underlying data stores.

SC1. The whole meaning of this service is that it should be publicly available. The confidentiality of the data relies upon some authorisation. If the asset is publicly exposed in the sense that the authorisation rules are bypassed, consumers with malicious intentions may get hold of the private data of other users. [$M \times M = M$]

SC2. Equivalent to SC1.

SC3. Equivalent to SC1.

SC4. The expected result of the process is to produce receipt data from underlying data stores and pass these to the web UI. If this web service fails to do so, it would mean that users can not find their receipts and will not be able to make guarantee or insurance claims on the purchased item. This would undoubtedly lead to trust issues between the user and the dSafe service. [$L \times M = L$]

SC6. Equivalent to SC4.

PA3: Receipt Parser

The worker process converts raw, incoming receipts to a dSafe-formatted normalised receipt. It also stores raw receipt files into a *original receipt* storage. The normalised receipt is stored and sent to another process for further work.

SC1. Making the receipt parser available would give external parties insight in the formatting of the raw receipts and how they are parsed. It could reveal any programmatic weaknesses that could be exploited in a directed DoS-attack, but this does not seem very likely. [$L \times L = L$]

SC2. Equivalent to SC1.

SC3. If manipulated, the output of the receipt parser could be altered, cascading an incorrect receipt into the *normalised receipt*-storage. Also, changes could be made in the *original receipt*-storage, which would make the discrepancy difficult to discover. [$L \times H = M$]

SC4. Expected results from this process are either normalised receipts being sent to storage and further processing, or the raw receipt being sent to the *original receipt*-storage. Unexpected results could be different information in the data outputs, between raw and processed receipts, the loss of either output, or the incorrect interpretation of input format or validity of raw receipts. [$M \times L = L$]

SC6. If unavailable, receipts would stack up in the queue. This would lead to a delay in the receipt storage and further processing, and failed receipt handling. [$M \times L = L$]

PA4: User Verification Service

This is a web service for allowing the stores to check whether a credit card is owned by a user of the dSafe application before sending the receipt to dSafe.

SC1. If the asset were to become available for all to use, the only information available would be whether or not a card number is owned by a user in the dSafe system. This can be exploited to check whether a credit card is valid. [$L{\times}M = L$]

SC2. An employee of the cloud service could check whether a credit card belongs to a user, or to check if the credit card is valid at all. [$L{\times}M = L$]

SC3. If an outsider could alter the web service, he/she could allow unwanted credit cards to be accepted as valid or make the service reject all valid request, telling the store not to send valid receipts. [$L{\times}M = L$]

SC4. Equivalent to SC3.

SC6. An absence of this service would mean that no receipts are sent to the application, leaving a set of purchases without electronic proof of purchase. [$L{\times}L = L$]

PA5: Contract Signing Service

This is the future module for signing contracts and other legally binding documents through the web interface. It should also be responsible for storing these documents, making them available to the correct users.

SC1. To sign a contract, one will need a high level security electronic signature, legally associated to a person. Without such a signature, the algorithm for signing the document should not be able to write a legally valid signature on the document. [$L{\times}L = L$]

SC2. To sign a contract, one will need a high level security electronic signature, legally associated to a person. Without such a signature, the algorithm for signing the document should not be able to write a legally valid signature on the document. Even though an employee could actually get a hold of the code behind the executing signer, this should not affect the validity of the documents signed. [$L{\times}L = L$]

SC3. One could have manipulated the function to sign a document with an invalid or no signature. This could lead the other party to believe that the document was signed. One could easily discover such a discrepancy. [$L{\times}M = L$]

SC4. Equivalent to SC3.

SC6. If the asset were unavailable, one could not sign or retrieve signed documents. This would make the whole signing/contract-service unavailable for the period of time, but the signing or retrieval process could be resumed. [$L{\times}L = L$]

Table 4.3: This table summarises the probabilities, consequences and resulting risk values for the data assets for each scenario.

Scenario	**Probability**	**Consequence**	**Risk**
Asset DA1 - End User Information			
SC1	L	L	L
SC2	L	L	L
SC3	L	M	L
SC5	L	M	L
SC6	L	L	L
Asset DA2 - Receipt Breakdown Information			
SC1	L	L	L
SC2	L	L	L
SC3	L	M	L
SC5	L	H	M
SC6	L	H	M
Asset DA3 - Unprocessed Receipts			
SC1	L	L	L
SC2	L	L	L
SC3	L	H	M
SC5	L	H	M
SC6	L	H	M
Asset DA4 - Original Receipt Storage			
SC1	L	M	L
SC2	L	L	L
SC3	L	L	L
SC5	L	M	L
SC6	L	H	M

SC1 – Widely published
SC2 – Employee gains access
SC3 – Manipulated by outsider
SC5 – Data changed
SC6 – Asset unavailable

Table 4.3: This table summarises the probabilities, consequences and resulting risk values for the data assets for each scenario. (continued)

Scenario	**Probability**	**Consequence**	**Risk**
Asset DA6 - Receipt Pointer Queue			
SC1	L	L	L
SC2	L	L	L
SC3	L	H	M
SC5	L	H	M
SC6	L	H	M
Asset DA7 - Contract Storage			
SC1	L	H	M
SC2	L	L	L
SC3	L	H	M
SC5	L	H	M
SC6	L	M	L
Asset DA5 - Failed Receipt Storage			
SC1	L	M	L
SC2	L	L	L
SC3	L	H	M
SC5	L	H	M
SC6	L	L	L

SC1 – Widely published

SC2 – Employee gains access

SC3 – Manipulated by outsider

SC5 – Data changed

SC6 – Asset unavailable

Table 4.4: This table summarises the probabilities, consequences and resulting risk values for the process assets for each scenario.

Scenario	**Probability**	**Consequence**	**Risk**
Asset PA1 - Receipt Upload Service			
SC1	M	H	H
SC2	L	M	L
SC3	L	H	M
SC4	L	H	M
SC6	M	H	H
Asset PA2 - Data Retrieval Service			
SC1	M	M	M
SC2	M	M	M
SC3	M	M	M
SC4	L	M	L
SC6	L	M	L
Asset PA3 - Receipt Parser			
SC1	L	L	L
SC2	L	L	L
SC3	L	H	M
SC4	M	L	L
SC6	M	L	L

SC1 – Widely published

SC2 – Employee gains access

SC3 – Manipulated by outsider

SC4 – Yields unexpected results

SC6 – Asset unavailable

Table 4.4: This table summarises the probabilities, consequences and resulting risk values for the process assets for each scenario. (continued)

Scenario	**Probability**	**Consequence**	**Risk**
Asset PA4 - User Verification Service			
SC1	L	M	L
SC2	L	M	L
SC3	L	M	L
SC4	L	M	L
SC6	L	L	L
Asset PA5 - Contract Signing Service			
SC1	L	L	L
SC2	L	L	L
SC3	L	M	L
SC4	L	M	L
SC6	L	L	L

SC1 – Widely published

SC2 – Employee gains access

SC3 – Manipulated by outsider

SC4 – Yields unexpected results

SC6 – Asset unavailable

Table 4.5: Listing of all assets with determinations as to which deployment models are acceptable

Asset	**Public**	**Private (internal)**	**Private (external)**	**Community**	**Hybrid**
DA1	✓	✓	✓	✓	✓
DA2	✓	✓	✓	✓	✓
DA3	✓	✓	✓	✓	✓
DA4	✓	✓	✓	✓	✓
DA5	✓	✓	✓	✓	✓
DA6	✓	✓	✓	✓	✓
DA7	✓	✓	✓	✓	✓
PA1	✓				✓
PA2	✓				✓
PA3	✓	✓	✓	✓	✓
PA4	✓				✓
PA5	✓	✓	✓	✓	✓

4.2 Deployment Model Acceptance

Having evaluated the risks for each asset we now need to determine which deployment models dSafe is prepared to accept for them. In Section 2.2 these models were described.

- **Public:** dSafe's initial choice of using Windows Azure as their cloud provider along with their rationale was stated in Section 3.2. Since Windows Azure is a public cloud this implies that they have no problem with the public cloud. Also, all assets fits into this model.
- **Private (internal):** This collides with dSafe's intention to avoid hardware and licensing costs as well as utilising the flexibility of the cloud and as such it is unacceptable. In addition, the web services (PA1, PA2 and PA4) can not be realised in this model.
- **Private (external):** This would give dSafe the benefits of hardware and licensing cost avoidance, but since the web services necessarily need to be exposed this model is inappropriate.
- **Community:** Even though dSafe could run on an inter-business community cloud with their business partners, this model would not be suitable for the current business model where users can access the information from outside such a community.
- **Hybrid:** dSafe has stated that they intend to use only one provider for their core services since the legal hurdles associated with one alone are non-trivial. This model can accomodate the deployment of all assets.

In Table 4.5 we summarise which deployment models have been deemed acceptable to dSafe.

4.3 The Twelve Domains of Critical Focus

In this section we will examine twelve domains that are relevant for both cloud and non-cloud systems (also shown in Figure 4.2):

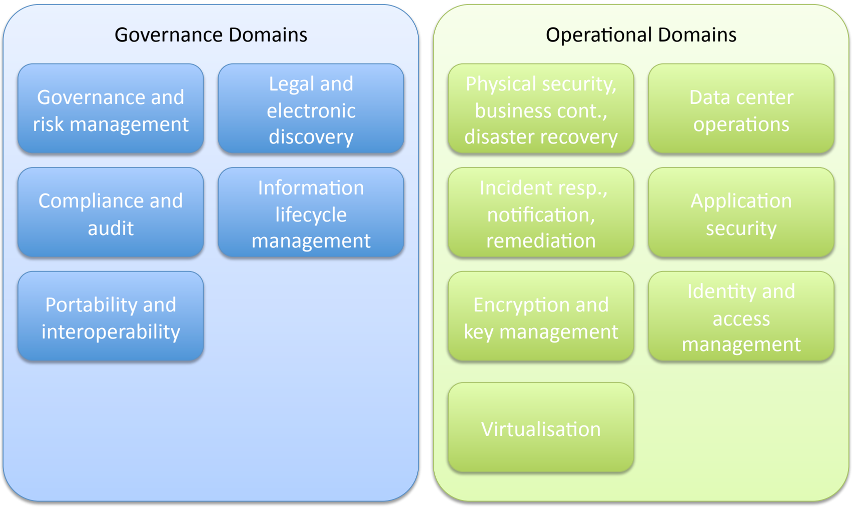
Governance Domains
Governance and risk management
Legal and electronic discovery
Compliance and audit
Information lifecycle management
Portability and interoperability
Operational Domains
Physical security, business cont., disaster recovery
Data center operations
Incident resp., notification, remediation
Application security
Encryption and key management
Identity and access management
Virtualisation

Figure 4.2: The twelve domains of critical focus

Governance domains:

1. Governance and risk management
2. Legal and electronic discovery
3. Compliance and audit
4. Information lifecycle management
5. Portability and interoperability

Operational domains:

6. Traditional security, business continuity and disaster recovery
7. Data center operations
8. Incident response, notification and remediation
9. Application security
10. Encryption and key management
11. Identity and access management

12. Virtualization

See [3] for further details on each domain. We introduce each domain and put it into the dSafe context. Each domain has important implications for dSafe that are inherent when putting enterprise systems into the cloud and so we will make recommendations in each of the these domains; recommendations aimed at making dSafe's application optimally secure.

4.3.1 Governance domains

The governance domains concern strategies and policy decisions regarding how to minimise risk and create a secure system in the cloud. An overview of the domains is shown in Figure 4.3.

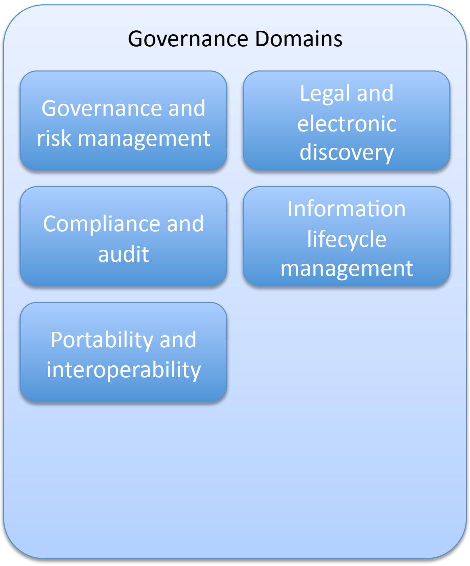
Governance Domains
Governance and risk management
Legal and electronic discovery
Compliance and audit
Information lifecycle management
Portability and interoperability

Figure 4.3: An overview of the five governance domains.

4.3.1.1 Governance and Risk Management

The Cloud Computing evangelists has for years told us about the hardware savings and the performance and flexibility gains achievable by moving services into the cloud. Unfortunately, companies deciding to try it out are not able to take the money and run. These monetary savings should be put into understanding and mitigating the many risks associated with allowing a cloud provider to store and manage your data. Governance and risk management is the ability to govern and measure these risks. This is done through thorough risk assessments of the application and the provider. A number of areas should be explored for a full understanding of the risks associated with migrating to the cloud.

> *Organizations should (. . .) assure reasonable information security across the information supply chain, encompassing providers and customers of Cloud Computing services and their supporting third party vendors, in any cloud deployment model.* – Cloud Security Alliance

Not much information exists out there of previous experiences and stories of the vendor meeting the customers' demands with regards to supporting a risk assessment directed towards the core processes of the vendor. As such a direct approach would seem implausible, if not impossible, for a small customer such as dSafe, the best way to look at how the vendor meets governance demands is to take a look at what is actually published with regards to how data centers are run and how the customer's demands are met. This would mean the openness of any security white papers, the concreteness of SLA, security process documentation, published success stories from other, third-party, customers or from audits done by respectable auditors with regards to established standards and certifications. Also, business continuity plans, incident management, contract requirements and other non-operational assets will be important to assess the complete risk of choosing a cloud provider. These items and more are covered in the preceding sections.

Recommendations: To ensure the validity of choice when establishing a contract with the cloud provider, dSafe should engage professionals for a review of contracts, SLA, any legal precedence, compliance with local law, End-User License Agreement (EULA) and terms of service in addition. A professional would be able to do a qualified risk assessment on the basis of these documents, and together with the following studies – both within governance- and operational domains, and the dSafe architectural risk assessment – it should state whether the choice is valid.

4.3.1.2 Legal and Electronic Discovery

dSafe intends to store citizens' personal data in another country. This has legal implications that need to be addressed, related to Norwegian as well as international law. The Norwegian Personal Data Act defines both regulatory and privacy requirements that dSafe needs to take seriously and so we will present a discussion on these requirements and make some recommendations in this section [44].

In Section 3.2.2 we detailed what types of information dSafe intends to store. The implications of storing this information prompted dSafe to contact the Norwegian Data Inspectorate to receive recommendations. The inspectorate's mandate includes:

> *(. . .) verifying that statutes and regulations which apply to the processing of personal data are complied with, and that errors or deficiencies are rectified.* – datatilsynet.no

However, dSafe are themselves responsible for making sure the relevant regulation is followed. The letter contained in Appendix A was the response to that inquiry. The relevant regulation is contained in the *Personal Data Act.* Section 8 of said act states:

> *Personal data (...) may only be processed if the data subject has consented thereto, or there is statutory authority for such processing, or the processing is necessary in order to fulfill a contract to which the data subject is party, or to take steps at the request of the data subject prior to entering into such a contract, (...)*

The act also defines the processing of sensitive personal data in the same manner, in Section 9, with the same provision for consent.

Section 13 clearly states that the security of the information must be maintained with regards to confidentiality, integrity, and accessibility in the connection of processing it. Simply put, this means that some sort of scheme must be employed to ensure that:

- Persons not involved with the processing does not gain access to the information.
- The data is not corrupted or changed en route.
- Persons given access are properly authorised.

Section 13 also states that the measures taken to ensure this, as well as the information system itself, must be well documented and made available to the controller, the data processor and the Data Inspectorate.

Section 15 has two clauses. The first being that data can not be manipulated in any other way than has been established in writing with the customer. The other being that data can not be turned over for storage or manipulation without such agreement. This becomes pertinent especially because dSafe intends to outsource storage to a cloud provider. The same agreement must include the measures that were demanded in Section 13.

Section 19 states that whenever data is collected from the data subjects themselves, the controller must inform the data subject of the identity of the controller, why the data is collected, whether the data will be disclosed and if so, the identity of the recipient and that supplying this data is voluntary.

While all these sections must be heeded, the Data Inspectorate emphasises Section 29 as dSafe's biggest challenge. The Section primarily states:

> *Personal data may only be transferred to countries which ensure an adequate level of protection of the data. Countries which have implemented Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data meet the requirement as regards an adequate level of protection.*

As with Section 15, this becomes pertinent due to dSafe's intention of outsourcing storage to a cloud provider. If the cloud provider then moves the data to another region which does not satisfy Section 29 has dSafe become in breach of Norwegian law? Section 30 defines a loophole:

> *Personal data may also be transferred to countries which do not ensure an adequate level of protection if the data subject has consented to the transfer.*

Here we again see that informed consent is very powerful. If this consent is attained from dSafe's users, then dSafe can be given considerable leeway in how the data are maintained and processed. Even if this condition has not been met, the Data Inspectorate may allow said transfer if dSafe provides adequate safeguards with respect to the rights of the data subject.

However, leveraging this informed consent can turn out to be unnecessary if the right cloud provider is chosen:

> *Microsoft must be able to comply with a myriad of regulatory obligations because it has data centers in a number of countries and offers online services to a global customer base.* – Microsoft*[29]*

Microsoft's Azure platform allows the setting of geolocalisation, which is the express guarantee of data staying in specific data centers. This would allow dSafe to adhere to the requirement of staying within the borders of a country that has implemented EU Directive 95/46/EC and thereby Section 29 of the Personal Data Act. Amazon also offers this feature under a different name whereas Google has no such feature. For more information on these features consult Section 4.3.2.2.

Recommendations: dSafe will need to create a clear and readable EULA to ensure a legally valid written consent is present. The EULA has to not only describe the relationship between the user and dSafe, but also the third party cloud provider. dSafe should aim to keep this as simple as possible by choosing this provider with due care. We readily recommend that dSafe utilises a cloud provider that offers the geolocalisation-feature as this will facilitate simpler conditions for the EULA. dSafe will also need to implement as well as document all features that ensure the security of all the information that flows through their systems to comply with Section 13 of the law. Refer to Section 4.3.1.3 for how best to achieve this compliance.

4.3.1.3 Compliance and Audit

In order to provide a meaningful EULA dSafe will have to get assurances from their chosen cloud provider on how data will be managed. If external auditors were to audit dSafe it is paramount that dSafe understands the capabilities of their chosen provider. The provider might not be able to offer any sort of guarantees. This could potentially be disruptive to certification processes and ultimately negatively affect dSafe's trustworthiness. The CSA mention several specific items dSafe must understand:

- Regulatory applicability for the chosen provider.

- The division of compliance responsibilities between dSafe and this provider.
- The cloud provider's ability to provide evidence needed for compliance.
- How dSafe is responsible for bridging the gap between the auditor and the provider.

ISO 27001 is becoming increasingly important and has been adopted by the ISO Norway (Norsk Standard). Several Norwegian business exist that can assist and/or certify dSafe as ISO 27001-compliant, e.g. *Center Teknologisk institutt Sertifisering AS*, *Det Norske Veritas* or *Nemko*. After contacting the Norwegian Data Inspectorate we were informed that they view the ISO 27001 process favourably as witnessed in this translated quote (Original correspondence in Appendix B):

> *We encourage using it [becoming ISO 27001 certified] if the company deems it useful with regards to their size and other needs. It will surely assist in implementing internal controls within the company with regards to employee training, discrepancy handling and auditing.* – Frank U. Eriksen, The Norwegian Data Inspectorate

Recommendations: In addition to formalising dSafe's information security policies this would have satisfied Section 13 of the Norwegian Personal Data Act (see Section 4.3.1.2). dSafe intends to put the entire system into the cloud so they will have to audit their chosen provider to ensure the correct management of data. This may prove to be difficult and implausible. If dSafe were to use a cloud provider that was ISO 27001 certified it would become easier for dSafe to bridge the gap between their auditor and their provider alleviating the need for dSafe to pursue an independent audit.

4.3.1.4 Information Lifecycle Management

dSafe intends to store and process large amounts of data, data that we have already shown to be potentially sensitive. This domain is about maintaining control over any and all data that flow through the systems. This becomes even more important when moving data into the cloud where a provider has another agenda: improving the bottom line for their shareholders.

The Cloud Security Alliance describes a version of Information Lifecycle Management: The Data Security Lifecycle. The cycle consists of six stages: Create, Store, Use, Share, Archive, Destroy. The CSA's keywords are listed here:

1. **Create:** *Classify* and *Assign Rights.*
2. **Store:** *Access Controls, Encryption, Rights Management, Content Discovery.*
3. **Use:** *Activity Monitoring and Enforcement, Rights management, Logical Controls, Application Security.*
4. **Share:** *Data Loss Prevention (DLP), Encryption, Logical controls, Application Security.*

5. **Archive:** *Encryption, Asset Management.*
6. **Destroy:** *Crypto-Shredding, Secure Deletion, Content Discovery.*

In the lifecycle, nothing is left to chance. Every scenario pertaining to data processing is covered.

We see that encryption is part of several of the steps in the cycle, but we will handle that subject under the operational domain of *Encryption and Key Management* in Section 4.3.2.5.

Recommendations: As the data moves from stage to stage in the Data Security Lifecycle, dSafe should maintain strict control of the data. Unless proper data security measures such as encryption and authentication are employed, dSafe will not be guaranteed that the data can not be sniffed by outsiders. The chosen provider should already have taken this into consideration by promoting the use of encryption.

dSafe should always be in control of where their data are physically. We have already seen that both Amazon and Microsoft offers this guarantee through allowing the express setting of data locality. This guarantee is essential due to the legal implications described in Section 4.3.1.2.

Optimally, from purely a security point of view, the chosen provider should have put into place facilities for offering encryption of data at rest in their storage units, at file system level. None of the cloud providers we surveyed has such facilities at the time or writing.

In cases where illegalities might have taken place in the data center, law enforcement agencies might target hardware in trying to get the data contained within. dSafe should ascertain assurances that their data is not seized along with data targeted by law enforcement agencies. This scenario could render dSafe's services unavailable merely because their services happened to be co-localised with another company.

dSafe needs to develop plans for backup and recovery. Many cloud providers offer a redundant storage solution that can guarantee the availability of dSafe's data, but dSafe will need to backup their data in case something unforeseen happens. For example: dSafe rolls out a new version of a part of their system that has not been sufficiently tested and the system starts losing data. No backup solution is complete without a recovery plan such that their services can come back online with correct data [47, Chapter 26]. In addition such a backup solution can help dSafe if they decide to migrate to another provider in the future.

The final stage of the Data Security Lifecycle is also highly dependent on the chosen cloud provider. When dSafe's policies mandate the deletion of data, they must have the provider's assurances that this data is not retained anywhere outside their control. In addition it is a well known fact that data on magnetic hard drives can be recovered to some degree depending on the mode of deletion. *Secure Deletion* implies that that data is made completely unrecoverable and dSafe will need to ascertain this guarantee from their chosen provider.

4.3.1.5 Portability and Interoperability

The business chooses a cloud provider and launches their services on the platform. Somewhere down the line many things could happen that would cause the the business to consider leaving this cloud provider. The CSA suggest several reasons for why this might be favourable: increased costs, provider going out of business, degradation in service quality or a dispute between business and vendor. Unless this situation was planned for, this could become costly. This is a non-trivial problem. The Cloud Security Alliance refers to this as *basic business continuity planning*. In this section we will deal with potential measures for easing continuity planning in a migratory setting. For the type of business continuity dealing with disaster recovery see Section 4.3.2.1.

- **Software-as-a-Service:** When dSafe's application is written to run like this the CSA finds it plausible that we are talking about rewriting the application for another provider. The focus should be to preserve or enhance the security functionality of the application.
- **Platform-as-a-Service:** In this case the focus remains the same, but it should be simpler since the application is more independent of the providers services.
- **Infrastructure-as-a-Service:** When dSafe manages the platform themselves, they will only need to find another IaaS provider. Even this is non-trivial as we will see next.

IaaS vendors utilise virtualisation technology. To ease migration it is important to understand the format of the virtual machine instances dSafe would be running on the current provider's services. For example Amazon describes that they utilise Xen Hypervisor[1] heavily modified and extended [20]. These details could prove disruptive if dSafe were to migrate the virtual machine away from AWS. The issues of virtualisation will be further elaborated upon in Section 4.3.2.7.

Recommendations: What happens when dSafe can not find a new provider that can accommodate their needs? dSafe should document and understand their needs sufficiently to be able to identify this situation. This includes, among other things, the sheer amount of data that needs to be migrated.

In Section 4.3.1.4 we recommended that dSafe pursued ISO 27001 certification. To become certified dSafe would have to have documented every security control thoroughly. If dSafe were to choose not to pursue this certification we here recommend that dSafe still document their architecture with focus on their security controls sufficiently. This will, in part, facilitate a successful migration with no loss of security.

The challenges of this domain depend to a large degree on which service model dSafe is working on, but regardless dSafe will need to have in place backup and recovery solutions to be able to successfully complete a migration. This was mentioned in

[1] See `http://xen.org/`

Section 4.3.1.4 but it is highly relevant here as well. Efficient backup and recovery procedures are vital in achieving maximum business continuity.

Another good practice for maintaining a system that is convenient to migrate is avoiding tying one's applications to the provider's proprietary APIs. A sound architecture could include abstraction layers to achieve this, albeit it is not within the scope of the current study to explore this field.

dSafe should commit itself to following as many open standards as possible. This will further facilitate migration.

4.3.2 Operational Domain

Whereas the governance domain concern broad security strategies the operational domains focuses on how to actually achieve the goals of these strategies. Hence we will be more specific and focus on what dSafe should actually do to keep their data secure in the cloud. An overview of the operational domains is shown in Figure 4.4.

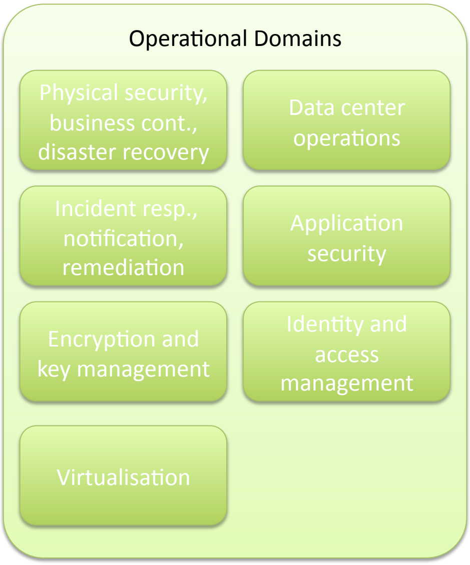
Operational Domains
Physical security, business cont., disaster recovery
Data center operations
Incident resp., notification, remediation
Application security
Encryption and key management
Identity and access management
Virtualisation

Figure 4.4: An overview of the seven operational domains.

4.3.2.1 Traditional Security, Business Continuity and Disaster Recovery

When companies start venturing out into the cloud this area grows more complex. dSafe might gain performance increases and cost savings by utilising the cloud, but the risks become more dire. *Business continuity* is the science of maintaining service functions on which businesses depend on. This is primarily about prevention techniques but *disaster recovery* is also part of it. Running several cloned services on different hardware is a typical example of the first whereas having an effective *backup and recovery*-solution is a good example of the latter. These issues becomes paramount for cloud providers for the simple fact that most providers provide subscription based pricing schemes; when services do not run the provider does not get paid.

In Section 4.1.2 we described and interpreted the six different scenarios from CSA's Security Guidance [3] of which the second deals with the risk of a rogue insider. We deem the probability to be low but it can be surprising what people will do when they believe they can get away with it. Consequently, this is a scenario all vendors bear in mind. In addition to audits and standardised verification, promises regarding data center security are also made on behalf of each provider through security whitepapers and articles. Little information was found on the policies of Google and Microsoft, and the facts found were related to data centers – not specifically AppEngine or Azure – so some assumptions are made in the above studies of these. Amazon states that no customer instances can be reached by Amazon's administrators. This is reassuring, but with regards to the facts present about the virtualisation technology run by other providers (see Section 2.4.10 for details), this may – although not explicitly stated – be the case for multiple other providers. If such facts were presented on behalf of the respective cloud technologies, this would be a crucial detail in the cloud provider decision-making process.

In Section 2.4.1 we saw that every vendor uses redundancy and geographic dispersement to ensure the security of the customers' data and applications. While Amazon disperses the data across different data centers, we could not locate any information on how the data is treated within the single data center. Microsoft, however, is very open about its *fault domains* within each data center.

Recommendations: To limit the possibility of having provider employees tampering with or gathering your data in the cloud, it is advisable to choose a cloud provider with defined data center access routines, a process for distributing privileges to administrators and a process for selecting employees which are to have physical access to the hardware where the data resides. These processes should be verified by some internal third party audit[2] and top-level routines publicly announced. All three cloud vendors have SAS 70 Type II verified data centers. This fulfills, to some degree, the recommendation of verified processes. In addition, Microsoft is ISO 27001 certified, which is an additional international verification (see Section 2.4.3 for further details). It is our opinion that dSafe should choose an ISO 27001 certified vendor so they would be able to look the vendor in the cards by requesting insight into the certification process.

The *fault domains* of the Windows Azure platform, to us, look very promising in ensuring both the security of the data and the continual availability of dSafe's services, making Azure look like a sound choice.

[2] An internal third-party audit is an external verification- or standardisation agency performing audit on routines and processes on behalf of the organisation in question.

4.3.2.2 Data Center Operations

To fully grasp the basic architectures and technological possibilities, the cloud customer should understand how the provider implements *The Essential Characteristics of Cloud Computing* – five characteristics demonstrating their relation to and separating cloud computing from traditional computing. The reader is referred to [3] for further descriptions of the characteristics.

On-demand self-service is the first characteristic. This is the ability to provide auto-scaling of the compute and storage cloud depending on instance load. With potential increasing traffic and little traffic history to draw conclusions from, it is essential for dSafe to employ any auto-scaling capabilities available in or supported by the cloud instance, as long as the costs associated are reasonable. This characteristic also addresses the *Rapid elasticity* and *Measured service* characteristics.

As described in Section 2.4.6, the providers differ somewhat in the way that this is handled. This is mostly due to the nature of the different clouds. AWS and Azure provide API for administering instances, whilst Google does not provide a notion of an *instance* at all. The solutions from Amazon and Windows are quite similar and no major differences separate the two. With Google's model, there is no need to administer the scale of the instance, since this is completely transparent to the user. From a technical perspective, dSafe is dependant on having a completely scalable design to meet the possible increased load in the future. To ensure that the hardware "leased" follow the meeting demand of processing, an API for the statically defined instances is necessary. If the instance is transparent, or at least dynamic, a comprehensive overview is needed – as the one provided by Google. In a business perspective, dSafe would benefit from having full control over the instances and therefore the cost model – although this is regarded as outside the scope of this study.

Broad network access *means to what degree is computational and storage capabilities available from different platforms?* As more and more internet services use XML-based, standardised protocols – available over HTTP – the services become more and more accessible from a host of different devices. *Broad network access* only encompasses HTTP-based protocols, since all dSafe services are accessible through web services, over HTTP or HTTPS. Both Windows Azure, Amazon Web Services and Google AppEngine support RESTful services, SOAP and other XML-based protocols and no relevant difference between the providers exist for this particular characteristic.

Resource pooling means grouping together resources, e.g. storage, processing, memory, network bandwidth and virtual machines. A cloud provider may specify the location of these resources for dSafe to be able to adhere to legal restraints (see Section 4.3.1.2 for further details) on the data handled by the hosted application.

Even though no promises are made for data locality, one can in both Windows Azure and Amazon Web Services define in which area the application should reside. This is essential for dSafe, planning to host all data – including potentially sensitive data – in the cloud. It seems that this locality mechanism is a strong selling point both for Microsoft and Amazon, so one could trust them to not move data outside this defined area. Google does not offer dSafe a location-specific environment.

Recommendations: For the *on-demand self-service* cloud characteristic, the solutions provided by each vendor is suitable for their cloud product model, although to prevent unwanted scaling of the dSafe instance, following from a high load (e.g. directed DoS), some control of the instance is preferable, provided by Microsoft and Amazon through manual, user-defined instance scaling, and by Google through the definition of an upper limit of cost for each application. In dSafe's context, all these three vendors would provide enough control of the instance.

With regards to the *broad network access* characteristic, all of the surveyed cloud vendors provide support for all relevant dSafe-employed protocols. There is therefore no decisive difference between vendors regarding this characteristic.

How resources are being allocated is not really of interest to dSafe. The interesting point is where dSafe's own, potentially sensitive, data are being processed and stored. Since all vendors should be trusted when it comes to efficient use of resource pooling, some will spread the data over wast geographical areas, whilst other will try to keep the data within a defined area. Both Microsoft and Amazon could be trusted to host the data in specific locations. In dSafe's context, Google does not provide sufficient promises of data locality.

4.3.2.3 Incident Response, Notification and Remediation

The next challenge dSafe has to face is what to do when security incidents, data breaches and such occur in the cloud. The CSA suggests augmenting existing standards for security incident response, bringing up the standards to the current reality where the responsibilities are shared between customer and cloud provider. Introducing measures to ensure security into the software's architecture plans from the beginning is a relatively new tradition. Hence most older applications are not prepared for being hosted off-site, much less designed for it. The usual challenges of local data center operations are complicated further by the nature of the public cloud: Co-tenancy with other the applications and data of other customers means making sure your data is not accessible to them. These co-tenants can even be malicious. In this section we will analyse the capabilities and strategies of the different providers in this field.

The CSA recommends that cloud providers consider application layer logging to avoid the clutter of mixing the log output for several customers' applications. We saw in Section 2.4.5 that every cloud vendor surveyed have implemented this and how this could be leveraged through the use of Security Information and Event Management (SIEM) systems.

Recommendations: In dSafe's current design documents, the designers have not included incident analysis or prevention systems. Regardless of which platform dSafe chooses their developers should implement or even buy an existing SIEM solution. Evaluating existing SIEM solutions falls outside the scope of this report but what is relevant is that dSafe understands the cloud providers' capabilities within logging and diagnostics to enable sufficient data mining and the creation of appropriate rulesets. In this section we have looked at the different logging capabilities of cloud providers and it is clear that both Microsoft and Amazon takes this very seriously, both publishing APIs to enable SIEM processes. For further information on what SIEM is and the how and when to leverage SIEM dSafe should consult [48].

4.3.2.4 Application Security

This domain is about deciding which service model is the most appropriate. The CSA likens the security requirements of an application residing in the cloud with the application residing in a DMZ. This in effect means designing defence into every layer of the application and expecting the very worst malicious hackers can throw at it. The domain covers security in the deployment phase, in the production phase and also the decommissioning phase. The CSA divide the consequences of cloud computing into five aspects:

- **Application Security Architecture:** The application in the cloud can have a very dynamic relationship with various third party dependencies.
- **Software Development Life Cycle:** The idea of including security into the software design process at the very beginning has steadily been gaining traction over the last decades. The introduction of cloud computing makes this even more important: Security implications can be severe and the "overhead" of implementing the needed security may make project leaders think again when considering cloud deployment.
- **Compliance:** The implementation of SIEM processes has very real consequences for security. Implementation is affected by the simple fact the the cloud system is a distributed system, creating a need for integrity and confidentiality in the way logs and incident reports are retrieved and processed. The nature of SIEM processes is discussed in Section 4.3.2.3.
- **Tools and Services:** Developing for the cloud might confer constraints on software libraries and developer tools that are available to the customer. This needs to be properly understood for the customer to be able to make proper choices.
- **Vulnerabilities:** The distributed nature of applications in the cloud brings with it new vulnerabilities related to communications between machines.

For IaaS where the customer creates their own virtual machines complete with guest operating systems the normal hardening techniques for machines residing in DMZs should be applied. AWS supplies all virtual machine instances with a firewall with a default of Deny All. Microsoft does not provide virtual machines with administrative access in the same way that Amazon does, going for a another approach. The web roles and worker roles of the Azure platform are designed to be web services so only ports 80 (HTTP) and 443 (HyperText Transfer Protocol Secure (HTTPS)) are available. No other incoming ports can be opened to the outside world. Applications running on Google's PaaS platform do not need controllable firewalls. The platform even locks down the use of sockets completely, only allowing HTTP or HTTPS requests to other hosts.

The CSA makes the explicit recommendation of always encrypting network streams. We will handle this subject under *Encryption and Key Management*, Section 4.3.2.5.

As another application security enhancing technique, the CSA recommends logging. The implementation and usage of Security Information and Event Management was discussed in Section 4.3.2.3.

Recommendations: dSafe should contractually make sure that they are allowed to perform remote vulnerability assessments such as penetration testing on their applications when they are deployed. Cloud providers may prove reluctant or hesitant to allow this since it can be difficult to seperate these attacks from real attacks. Nevertheless, it is extremely important that dSafe is able to do vulnerability assessments on their production environment.

It is paramount that dSafe properly understands the security implications inherent when designing an application for the cloud. The CSA emphasises three items within the generalised software development life cycle that are affected by the nature of the cloud:

1. **Threat and Trust Models:** Bringing in third party providers complicates these models further. Earlier domains have explored this in further detail. E.g. certifications and other audits can mitigate threats and build trust.
2. **Application Assessment Tools:** The different cloud providers have different APIs that leverage their specific platform's features as well as development tools. We mentioned it in the domain of *Interoperability* that API abstraction can be used to counteract this.
3. **Software Development LifeCycle (SDLC) and Quality Assurance (QA) Processes:** For all intents and purposes, the public cloud is like the DMZ. Hence, dSafe should protect every asset in their system as if they were residing in a DMZ. This means that no assets should be accessible without proper authenticated credentials. *Identity and Access Management* is covered in Section 4.3.2.6.

4.3.2.5 Encryption and Key Management

We have described that dSafe intends to store personal data and potentially sensitive personal data depending on the end-users. In Section 4.3.1.2 we saw that Norwegian law demands that information is inaccessible to untoward people and is protected in transit from both exposure and being altered. Encryption assists greatly in achieving compliance within these very important problems of confidentiality and integrity. The three decisions that need to be made are: Should data be encrypted while:

- In transit?
- At rest?
- On backup media?

Furthermore, the challenges of the management of encryption keys are non-trivial. The CSA defines the following three areas of focus:

- **Secure Key Stores:** As with the data these keys are used to encrypt and decrypt, the keys themselves needs to be protected to the same degree.
- **Access to Key Stores:** The access to the keys of course needs to be tightly controlled to keep all data from being decrypted and stolen.
- **Key Backup and Recoverability:** While the intentional loss of an encryption key could be said to be the most effective secure deletion method, the unintentional loss of an encryption key can cripple the company by rendering their stored data effectively lost.

SSL will encrypt end-to-end-traffic within the application layer. This will make it harder for attackers to get a hold of any authentication data transferred for access to the service. SSL encryption is available on all cloud platforms, although the requirements and implementation differ. On the Windows Azure platform developers are able to upload SSL certificates through a web portal and associate them with services they wish to provide. Every component in the AppFabric is available as a web service. This certificate process itself is performed over an SSL protected connection to prevent eavesdropping [49]. In addition, Microsoft enforces the use of SSL whenever connecting to SQL Azure [50]. Amazon highly recommends SSL connections for everything. Beyond this, Amazon does not provide facilities to assist the client in implementing it. On Google's AppEngine SSL must be enabled for the Uniform Resource Locators (URLs) through which you wish to support secure traffic. If the application is served outside a Google Apps domain, you must direct all traffic though the application's appspot domain.

None of the inspected cloud providers provide facilities for encrypting any data in any of the surveyed data storage solutions. Microsoft's SQL Server 2008 R2 solution supports what they call *Transparent Data Encryption*[3], but this functionality has not made it into the cloud version. According to MSDN Magazine, however, it is a feature being considered for future releases [51].

Recommendations: To satisfy Norwegian law dSafe should encrypt data while in transit. SSL is a tried and true solution for this problem. We also recommend that dSafe encrypts any data at rest and on backup media. This will effectively mitigate the first and second scenario (See Section 4.1.2) which are about the

[3] See `http://msdn.microsoft.com/en-us/library/bb934049.aspx`

unintended disclosure of information. Unfortunately this will have to be implemented on the application level since no provider supports this. Provided dSafe chooses a proven algorithm and an appropriate key size the developers would also have solved the *Secure Deletion* problem described in Section 4.3.1.4. The keys themselves should not be kept in storage at the same provider since this would counteract the mitigation. Whatever the developers of dSafe choose to do they should in no circumstances use any keys provided by the cloud provider in their own cryptographic functions for the obvious reason that the provider would then have access to the information.

4.3.2.6 Identity and Access Management

Authentication is the process of identifying a peer before exchanging a piece of information. Many different standards and implementation for authentication, and the cloud providers seem to differ in their choice. If a standard is adopted by one of the large financial services partnered with dSafe, and federated authentication were to become available, this would be a deciding factor for mitigating any authentication risks and usability issues associated with a decentralised authentication scheme.

As the BankID[4] identification solution is becoming the de facto standard for banking Electronic ID (eID) in Norway [52], this – or a similar solution – would be preferable in order to already take advantage of a wide spread distribution of key generators and other items needed for a strong two-factor authentication. The security tiers, or levels, of Norwegian governmental eID solutions are defined by the *Framework for authentication and non-repudiation i electronic communication with and in public sector* [53], a framework published by the *Ministry of Government Administration, Reform and Church Affairs* to aid governmental operations in the process of securing collaboration in open or closed networks. We will also adhere to this norm while defining eID security levels. The framework defines four levels of security (translated from Norwegian):

Level 1 – No requirements on authentication:
Gives little or no security. Open solutions fall under this category. Some solutions fall under this category because they are unable to meet the requirements of level 2. Examples of solutions in this category are:

- Self-defined password and username online.
- Identification with Social Security Number (SSN).

Level 2 – One-factor authentication:
Extended requirements on authentication factor issuing, the security of persisted authentication credentials, non-repudiation requirements. Examples of solutions in this category are:

- Static password, generated and sent to a *directory of residents*-registered (referred to as *registered*) address.

[4] See `http://www.bankid.no`

- Password calculators without password protection, at least distributed to a registered address.
- *Single-use password*-list distributed to a registered address.

Level 3 – Two-factor, one dynamic:
Additional procedures for issuing the authentication credentials from level 2, and more strict security of stored credentials. Same non-repudiation requirements as for level 2 apply. Examples of solutions in this category are:

- Password generators protected with a PIN-code, where the first PIN is sent in a separate shipment.
- Single-use password sent to cellular phone, where the phone is registered with a code sent to registered address.
- *Person-Standard* certificate security level defined in *Requirements specification for PKI for the public sector* [54]. The reader is directed to the specification for further details.
- Single-use password-lists used together with defined password and username. The definition of the password should be done with the use of a single-use code sent to registered address.

Level 4 – Two-factor, one dynamic:
According to current regulations, the solutions must be self-declared in *The Norwegian Post and Telecommunications Authority (NPT)* in relation to their fulfillment of the requirements defined by the *Requirements specification for PKI for the public sector* regarding the categories *Person-High* and *Enterprise.* Examples of technologies that could meet the requirements of this security level are:

- Two-factor solution where one of the factors is dynamic, and where one of the authentication factors or a registry process factor is issued personally. A third party is used to register a log with the mapping between the information, or any actions performed on this, and the identity of the user. The log shall be stored with a protection against modifications.
- Two-factor solution where one of the factors is dynamic, and where one of the authentication factors or a registry process factor is issued personally. Special software to ensure complete non-repudiation of the logging and to hinder any operator modifications against such logs.

Furthermore, the NPT define several levels of certificate security in the *Requirements specification for PKI for the public sector.* These levels are *Person-High*, *Person-Standard* and *Enterprise* and is further described in the specification. Solutions that meet requirements for *Person-High* could satisfy the requirements for *level 4* in the framework.

A level 4 accepted eID (with a qualified[5] electronic signature) is also valid for level 3, 2 and 1; a level 3 accepted eID (with an advanced electronic signature) is a valid level 2 and 1 eID etc. This means that if a level 2 eID is required to log into a service, e.g. the dSafe receipt service, one can employ a level 4, 3 or 2 eID for this purpose.

Several of the eID-issuers already on the market in Norway are both defined as qualified and Public Key Infrastructure (PKI)-based with a *Person-High* or *Enterprise* security level. Examples of these are Buypass[6], BankID and Commfides[7]. Both Buypass and Commfides are regarded as level 4 solutions. BankID is by some sources [55, 56] considered as level 4, but did not get accepted as a valid eID for the governmental eID-portal ID-porten due to some issues with the lack of message encryption [57, 58, 59]. A number of scandals imposed by the lack of secure authentication, especially two-factor authentication, have been uncovered during the last couple of years [60] and the Citizens Financial Bank[8] even got sentenced in 2007 due to neglecting a document [61] recommending two-factor authorisation for banks.

The scope-creep defined in Section 4.1.1 encompasses the signing of documents through the dSafe service. Since *signing* of documents is now introduced, this would impose new security requirements beyond those that follow the receipt *reading* service. The Norwegian *E-Signature Law (ESL)*[62] facilitate a secure and efficient use of electronic signatures by defining requirements to qualified certificates, to the issuers and to safe generation of these signatures. We do not see why any of the level 4 candidates mentioned above could be used for electronic signatures according to the ESL, as we have no experience as legal practitioners. BankID's signing capability approach has been legally reviewed and found valid [63]. All three vendors are self-declared providers of legally valid electronic signatures [64, 65, 66] and BankID is already used for loan document signing in two Norwegian banks, DnBNOR and Postbanken.

Recommendations: Since one might reason about shopping patterns and whereabouts from the receipts stored in the dSafe data stores, one must protect these properly, with the right authentication level. The data stored by dSafe is not, per se, defined as *sensitive* data by Norwegian law (see 3.2.2 for further discussions). Since the data stored by dSafe in many ways are of the same nature as web based financial solutions, and in addition open for reasoning around the purchase data, one should – even though one can not make any financial transactions through the dSafe system – require the same authentication level as an internet banking application. Since all internet banking is recommended to use two-factor authentication – which is not a requirement on level 1 and 2 – one must employ an eID-solution with at least level 3 security.

One of the scope creeps identified is the electronic document signing functionality.

[5] A *qualified* electronic signature is an advanced electronic signature with additional requirements to the issuer of the eID and the signing equipment.

[6] See `http://www.buypass.no/`

[7] See `http://www.commfides.no/`

[8] See `http://www.citz.com/`

It will sign and store legal documents for the signees, with their personal signatures. dSafe should require a minimum level 3 eID for the receipt services. With regards to the signature service, one should choose an eID provider with proven security in accordance to Norwegian law (the ESL). At the time of implementation of this service, at least the above mentioned providers should be reevaluated to ensure adequate security and conformance to the levels set by established signing agencies and law. For now, it seems that BankID, Commfides and Buypass are all adequate solutions.

4.3.2.7 Virtualization

One of the major issues with cloud computing is the fact that the application and the data are residing in an instance in the same environment as thousands of other instances. What separates one instance from the other can be the use of virtualisation – to allow the user to see the instance as a separated hardware unit, without actually being physically separated. Due to the nature of the dSafe application, it is highly dependent on being separated from other data, at least logically. The in-depth study of the virtualisation technology for each vendor is presented in Section 2.4.10

Both Azure and AWS rely on a customised, commercial hypervisor for the virtual instances. Although somewhat different in architecture, the security they provide, and the measures taken are similar. The main visible security characteristic that separates the two are the more apparent possibility for the user to deploy third-party security software on the AWS instance. The Azure hypervisor seems to be more proprietary, both with regards to the closed-source hypervisor and the lack of appearance of non-Microsoft security software. What separates the Google AppEngine from the other vendors is the lack of hardware virtualisation, the type of virtualisation relevant for, and addressed in, this study.

The instance separation is similar on Windows Azure and Amazon Web Services, but the Google AppEngine application separation is not completely mapped out. The virtual machine environments are built on proven technology, so separation should be adequate. Google states to have included security measures in multiple tiers, but this is strictly disclosed [67].

One of the other areas discussed in the virtualisation study was the management interface security. The programmatic management interfaces (Azure and Amazon) are protected by an authentication requiring a X.509 certificate and is encrypted point-to-point by SSL. All web interfaces are protected by a one-factor password authentication over SSL. Since Google does not provide instance management, and therefore no management API, this yields no decisive difference regarding management interface security.

Recommendations: The dSafe application is dependent on being hosted on a completely compartmentalised environment in order to protect data. This is something that is addressed by all vendors, but not much is known about Google's approach. Amazon is the only vendor providing extensive insight into virtualisation implementation through the open-source hypervisor, therefore making this approach seem recommendable. Azure use a customised version Hyper-V, technology proven to be quite secure through extensive commercial use. The third-party support of AWS is best-in-class. Since management interface security seems to be similar for all three vendors, separation of instances and hypervisor proof-of-quality seems similar for Azure and AWS. dSafe is recommended to use either AWS or Azure. AppEngine may provide similar, or even better, security measures on some of the characteristics studied, but no concrete information about these implementations exist, so no conclusion is drawn with regards to AppEngine and "virtualisation".

4.4 Risk Mitigation

In this section mitigations are extracted from the preceding discussion. The following sections detail the reasoning for each domain while Tables 4.6, 4.7 and 4.8 summarise the capabilities of respectively Microsoft, Amazon and Google. The prefixes *GOV* and *OP* refer to *governance* and *operational* and are used to index the three tables. The scenarios were described in Section 4.1.2.

Governance and Risk Management (GOV1)

The tables list that, for all assets, the vendors somewhat mitigate scenario two through Terms of Service agreements that make guarantees as to how and when employees access customer systems and six through all the measures they take to be able to guarantee availability in their SLAs.

Legal and Electronic Discovery (GOV2)

Both Microsoft and Amazon mitigate the data asset scenarios one, three and five, by offering a geolocalisation mechanism; the express ability to limit where data is stored. By using the feature, the customer can partially mitigate the danger inherent when data is distributed. Google unfortunately offers no such control.

Compliance and Audit (GOV3)

All three vendors somewhat mitigate scenarios one, two and three for all assets and five for data assets by being SAS70 certified. Microsoft is ISO certified as well, providing an even higher commitment to being open to how information is handled and protected from unauthorised users.

Information Lifecycle Management (GOV4)

Microsoft's commitment to becoming ISO certified somewhat mitigates the two first scenarios for all data assets. The certification documents include, among other things, what happens with harddrives that are retired so dSafe can be assured that their data on the decommissioned disks is not distributed to anyone.

Portability and Interoperability (GOV5)

We were unable to find any mitigations for the scenarios described within this domain. While Microsoft, for instance, has an SQL Migration Tool to ease migration between platforms this does not increase the security of their platform.

Traditional Security, Business Continuity and Disaster Recovery (OP6)

All three vendors take their traditional security very seriously as we have previously described. E.g. Microsoft specifies that their data centers have armed guards protecting them [28].

Data Center Operations (OP7)

Again, both Microsoft and Amazon allow dSafe to control the data center in which its data and services shall reside. Microsoft's geolocalisation mechanism offers finer granularity than Amazon's equivalent feature. Therefore, Microsoft also somewhat mitigates scenario six which deals with preventing unavailability.

Incident Response, Notification and Remediation (OP8)

In this domain we have focused on the logging and auditing capabilities of the vendors. Both Microsoft and Amazon offer extensive API's for enabling the creation of SIEM systems (See Section 4.3.2.3).

Application Security (OP9)

None of the vendors are likely ever to allow their customers to do penetration testing on their production environment. E.g. this is due to the aspect of co-tenancy and thus understandable. However it means that none of the three vendors contribute any mitigation in this aspect.

Table 4.6: Microsoft Azure mitigations within the critical domains

Domain[e]	**Mitigates**[a]
Governance domains	
GOV1	*:SC2[b], *:SC6
GOV2	DA*:SC1[c], DA*:SC3, DA*:SC5
GOV3	*:SC1, *:SC2, *:SC3, DA*:SC5
GOV4	DA*:SC1, DA*:SC2,
GOV5	N/A
Operational domains	
OP6	*
OP7	*:SC5, *:SC6
OP8	*:SC2, *:SC3, *:SC4, *:SC5
OP9	N/A
OP10	DA*:SC1, DA*:SC3, PA1:SC1, PA2:SC1, PA4:SC1
OP11	N/A
OP12	*

a Annotated with asset and scenario: e.g. DA1:SC1.

b Valid for all assets for this/these scenario(s).

c Valid for all data assets for this/these scenario(s).

d Valid for all process assets this/these scenario(s).

e See Section 4.3

Encryption and Key Management (OP10)

Microsoft promotes SSL encryption by offering key management in Azure as well as having created easy to use interfaces to create SSL protected services. This contributes in mitigating eavesdropping (scenario one) and man-in-the-middle attacks (scenario three) for all data assets as well as scenario one for process assets PA1, PA2 and PA4.

Identity and Access Management (OP11)

None of the vendors provide authentication and identification facilities that are satisfactory to dSafe's needs.

Virtualisation (OP12)

All three vendors compartmentalise the applications running on their infrastructure. This contributes to mitigate every scenario for every asset. Exactly how much is a more difficult question.

Table 4.7: Amazon Web Services mitigations within the critical domains

Domain[e]	**Mitigates**[a]
Governance domains	
GOV1	*:SC2[b], *:SC6
GOV2	DA*:SC1[c], DA*:SC3, DA*:SC5
GOV3	*:SC1, *:SC2, *:SC3, DA*:SC5
GOV4	N/A
GOV5	N/A
Operational domains	
OP6	*
OP7	*:SC5
OP8	*:SC2, *:SC3, *:SC4, *:SC5
OP9	N/A
OP10	N/A
OP11	N/A
OP12	*

a Annotated with asset and scenario: e.g. DA1:SC1.

b Valid for all assets for this/these scenario(s).

c Valid for all data assets for this/these scenario(s).

d Valid for all process assets this/these scenario(s).

e See Section 4.3

Table 4.8: Google AppEngine mitigations within the critical domains

Domain[e]	**Mitigates**[a]
Governance domains	
GOV1	*:SC2[b], *:SC6
GOV2	N/A
GOV3	*:SC1, *:SC2, *:SC3, DA*:SC5
GOV4	N/A
GOV5	N/A
Operational domains	
OP6	*
OP7	N/A
OP8	N/A
OP9	N/A
OP10	N/A
OP11	N/A
OP12	*

[a] Annotated with asset and scenario: e.g. DA1:SC1.

[b] Valid for all assets for this/these scenario(s).

[c] Valid for all data assets for this/these scenario(s).

[d] Valid for all process assets this/these scenario(s).

[e] See Section 4.3

Chapter 5

Conclusions

As per our described approach we have been through a comparative study of cloud vendor platforms, followed by a thorough run-down of dSafe's challenges in the cloud. Consequently, it is time to answer the questions from Section 1.1:

- *Q1: Can information systems be as secure in the cloud as they would be in an on-premises environment?*
- *Q2: How can the planned services of dSafe be sufficiently secured in the cloud?*

5.1 Final Recommendations

In answering the second research questions we have made specific recommendations as to how dSafe should proceed. The first recommendation related to governmental operations is for dSafe to *hire legal aid to analyse license agreements.* dSafe needs to thoroughly understand the license agreements of the vendor they choose to ensure that dSafe knows what it's recourse is when security incidents occur. This is not intended to directly neutralise the threats posed in the scenarios but it will help dSafe in understanding any legal recourses that may be available if any scenarios should occur. dSafe should also *create an End-User License Agreement* on order to attain personal and informed consent from its users before storing their potentially sensitive data. To ensure the prevention of data loss, dSafe must design and implement procedures for performing backups and restoration of all their databases. This can be done through *creating a backup and recovery scheme.* dSafe should be acutely aware of how information flows through their system and, by proxy, the chosen vendor's system. *Documenting the information lifecycle* will contribute in mitigating the risks of vendor employees snooping (scenario two). dSafe should *choose open protocols and standards* to promote security and interoperability between cloud technologies, e.g. SAML and REST. With an ISO 27001-certified vendor, dSafe can request documentation from the thorough certification process, through a written application. This will help in understanding the vendor's information lifecycle and security measures, contributing to mitigate scenario two.

The nature of the data handled by the receipt storage requires, by law, that the data is kept within countries that implement EU Directive 95/46/EC. This requires *a vendor that allows explicit geolocation control.* Hence, among the surveyed vendors, Google is currently off-limits. To implement log analysis schemes so they can audit access on the cloud resident system, *a SIEM system should be employed.* This will facilitate further certification. dSafe needs to understand the magnitude of threats in the public cloud compared to the private one. Outsourcing data center operations enters a third party into the relationship that would otherwise only consist of two parties; dSafe and their users. Hopefully, this report has aided them in *updating the threat/trust model*, at least to the current situation. *All assets should be treated as if they are placed in a DMZ* and therefore protect all assets accordingly. This is imperative if dSafe's information is to be placed in a public cloud. This will indirectly help mitigate every scenario for every asset by assuming that all assets need to be protected. As required by Norwegian law dSafe must *protect data* in transit and ensure that only properly authenticated people are able to access their data. *Encryption* is an excellent way to achieve this. dSafe should *implement level 3 authentication* for their current design. When the scope creep occurs, dSafe must reevaluate their solution according to Norwegian law. Level 3 will effectively mitigate scenarios one and two. Google does not seem to use any established virtualisation technology in separating the applications of the customers to *ensure separation and compartementalisation*, whilst Microsoft and Amazon use proven commercial platform as a basis to achieve this – allowing dSafe to gain deeper insight into the validity of these.

Furthermore, in Section 4.4 we went through all the domains and described how the different vendors offer features that contribute in mitigating the risk scenarios. While there is no way to accurately score the different vendors based on the resulting tables it is our opinion that Microsoft seems to be the platform most appropriate to dSafe's needs.

Ultimately we find that dSafe's information system can become sufficiently secure if dSafe takes the new cloud-related risks into account. dSafe of course has to evaluate which of these recommendations they want to follow, and hopefully they will then be able to complete their information system in a secure manner.

5.2 Conclusions

It is our opinion, regarding the first research question, that systems can be built to be secure in the cloud, but as with on-premises deployments these are not necessarily *secure by default.* The security implications of the cloud have not yet become completely transparent to the developer. Creating information systems on cloud platforms requires even more diligence in security aspects compared to non-cloud deployments due to this new set of challenges. Customers of cloud computing services include high profile businesses such as Netflix (movie rental service on AWS), Siemens (software delivery service on Azure) and CapGemini (SaaS applications from Google Apps). In addition, security heavy-hitter VeriSign partner with Microsoft to enable encryption services on Windows Azure, thereby publishing a justification for the use of said platform in security critical applications.

As we stated in Section 2.5, in most aspects, the vendors' offerings are not significantly different. After going through all twelve security domains, however, we have found some important differences that lead to our conclusions. While Amazon and Microsoft have specifically targeted the enterprise customers Google have neglected them so far. The *AppEngine for Business*-program can change this in the future but Amazon has had a mature and proven platform for nearly a decade. Microsoft's Windows Azure, though only recently launched, is almost up to par with Amazon as far as their general feature sets go.

In Section 4.4 we summarised what risk scenarios the vendors mitigate in dSafe's case. While some of the reasons are specific to dSafe, there are also some that apply in the general case. We made the case for why dSafe should choose a vendor that is ISO 27001 certified. While we view this criteria as negotiable in dSafe's case there is one satisfactory vendor among the surveyed: Microsoft. Choosing Windows Azure would put customers into relations with a vendor that is open (enough to pursue this certification at least) about its information security practices and customers should use the documentation produced during Microsoft's certification process to supplement their own documentation. Another general mitigation is that Microsoft and Amazon both offer satisfactory APIs for logging needs and as such we do not recommend one over the other. Google, however, exposes much less information in their logging API.

Applications can become secure in the cloud, but this requires thorough risk analysis and assessments of both application and vendors – together with guidelines that apply to deploying to your chosen cloud platform. A framework designed specifically for the process of migrating applications to the cloud should be employed to ensure that all aspects are considered. What the cloud environment does not change is that applications are still no more secure than they are designed to be, regardless of residence, on-premises or in the cloud.

5.3 Further Work

Some areas were left unexplored in this study, due to the limitation of scope and time. This section will give some final pointers to areas in interest of dSafe and other companies migrating to the cloud.

5.3.1 Comprehensible Risk Analysis

As stated previously, the framework followed as a part of this study did not yield a complete risk analysis, but a rather shallow one, directed at the risks of choosing one or the other cloud vendor with regards to the major assets of the application. A complete risk analysis would not only encompass clear security issues, but also business, judicial, and strategic risks associated with having an application residing in the cloud. These analyses remain to be done.

5.3.2 Federated Identity

In Section 4.3.2.6, we discussed authentication with regards to the dSafe application in the cloud. Some recommendations were given with respect to security levels of receipt storage authentication and any future signing services. Different solutions already exist to assist such an implementation, and federated secure identities are currently available from companies such as Signicat. A study should be done to map which solutions seem feasible with regards to security in implementation and cost of use. The chosen solution should adhere to the recommendations given regarding the security levels previously described and discussed.

5.3.3 Maturity of Cloud Technology

In this report, three cloud computing vendors are studied, but the concept of *cloud computing* is not generally explored. To better understand the future development of cloud computing, an effort should be put into understanding the future of this model, how to meet the demands of the future, and how to grow securely in order to meet tomorrow's even more demanding security threats.

References

[1] B. Schneier, “Be careful when you come to put your trust in the clouds.” http://www.guardian.co.uk/technology/2009/jun/04/bruce-schneier-cloud-computing, June 2009. Accessed on March 21st, 2010.

[2] D. Talbot, “Vulnerability Seen in Amazon’s Cloud-Computing.” http://www.technologyreview.com/computing/23792/, October 2009. Accessed on March 21st, 2010.

[3] Cloud Security Alliance, “Security Guidance for Critical Areas of Focus in Cloud Computing V2.1.” http://www.cloudsecurityalliance.org/csaguide.pdf, December 2009. Accessed on March 22nd, 2010.

[4] L. Youseff and D. Butrico M., Da Silva, “Toward a Unified Ontology of Cloud Computing,” 2005.

[5] K. Rangan, A. Cooke, and M. Dhruv, “The Cloud Wars: $100+ billion at stake,” tech. rep., Merrill Lynch, 2008.

[6] P. Mell and T. Grance, “The NIST Definition of Cloud Computing version 15.” http://csrc.nist.gov/groups/SNS/cloud-computing, July 2009. Accessed on April 19th, 2010.

[7] D. Chappell, “Introducing Windows Azure.” http://download.microsoft.com/download/0/8/7/087A3AE1-2880-4452-88DD-09398D0A522A/Introducing_Windows_Azure.doc, Nov 2009. Last accessed on March 15th, 2009.

[8] K. Gibbs, “Campfire One: Introducing Google App Engine (pt. 3).” http://www.youtube.com/watch?v=oG6Ac7d-Nx8, 2008. Google Code video.

[9] Google, “Google Storage for Developers.” http://code.google.com/intl/no-NO/apis/storage/, 2010. Accessed on May 21st, 2010.

[10] J. Brodkin, “Gartner: Seven cloud-computing security risks.” http://www.infoworld.com/d/security-central/gartner-seven-cloud-computing-security-risks-853, July 2008. Accessed on March 22nd, 2010.

[11] Jericho Forum, “Jericho Forum Commandments.” http://www.opengroup.org/jericho/commandments_v1.2.pdf, 2007. Accessed on May 28th, 2010.

[12] Microsoft Corporation, "Azure Compute SLA v1.4." http://www.microsoft.com/windowsazure/sla/, April 2010. Accessed on May 3rd, 2010.

[13] Microsoft Corporation, "Azure Storage SLA v1.4." http://www.microsoft.com/windowsazure/sla/, April 2010. Accessed on May 3rd, 2010.

[14] Microsoft Corporation, "SQL Azure SLA v1.4." http://www.microsoft.com/windowsazure/sla/, April 2010. Accessed on April 19th, 2010.

[15] Microsoft Corporation, "AppFabric SLA v1.4." http://www.microsoft.com/windowsazure/sla/, April 2010. Accessed on May 3rd, 2010.

[16] H. Kommalapati, "Windows Azure Platform for Enterprises." MSDN Magazine - http://msdn.microsoft.com/en-us/magazine/ee309870.aspx, February 2010. Accessed on May 3rd, 2010.

[17] Amazon.com, Inc, "Amazon S3 Service Level Agreement." http://aws.amazon.com/s3-sla/, October 2007. Accessed on May 3rd, 2010.

[18] Amazon.com, Inc, "Amazon EC2 Service Level Agreement." http://aws.amazon.com/ec2-sla/, October 2008. Accessed on May 3rd, 2010.

[19] Amazon.com, Inc, "Amazon Elastic Compute Cloud - User Guide." http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/index.html, November 2009. Accessed on May 3rd, 2010.

[20] AWS Security Team, "Amazon Web Services: Overview of Security Processes." http://s3.amazonaws.com/aws_blog/AWS_Security_Whitepaper_2008_09.pdf, September 2008. Accessed on April 19th, 2010.

[21] Google, "App Engine for Business SLA (draft)." http://code.google.com/intl/no-NO/appengine/business/sla.html, May 2010. Accessed on May 22nd, 2010.

[22] Google, "What is Google App Engine?." http://code.google.com/appengine/docs/whatisgoogleappengine.html, 2010. Accessed on May 6th, 2010.

[23] IBM, "Web Services Federation Language." http://www.ibm.com/developerworks/library/specification/ws-fed/, 2007. Accessed on February 22nd, 2010.

[24] J. Bodley-Scott, "Access or Identity - Chicken or Egg?." https://www.opengroup.org/jericho/IDM2009_jbs.pdf, November 2009. Accessed on April 7th, 2010.

[25] K. Brown, "A Developer's Guide to Access Control in the Windows Azure platform AppFabric," November 2009. Accessed on February 22nd, 2010.

[26] M. Aller, "Banks and OAuth support." http://blog.maxaller.name/2009/08/banks-and-oauth-support/, August 2009. Accessed on May 3rd, 2010.

[27] ISM3 Consortium, "The Information Security Management Maturity Model v2.1." http://www.ism3.com/page1.php. Accessed on May 10th, 2010.

[28] Microsoft TechNet, "VIDEO: Real World Azure - Windows Azure Security." http://edge.technet.com/Media/Real-World-Azure-Windows-Azure-Security/, March 2010. Accessed on May 30th, 2010.

[29] Microsoft Corporation, “Securing Microsoft’s Cloud Infrastructure.” http://www.globalfoundationservices.com/security/documents/SecuringtheMSCloudMay09.pdf, May 2009. Accessed on April 7th, 2010.

[30] Microsoft Corporation, “Windows Azure Platform - Support.” http://www.microsoft.com/windowsazure/support/, 2010. Accessed on May 22nd, 2010.

[31] Amazon.com, Inc, “Amazon Web Services Web Premium Support.” http://aws.amazon.com/premiumsupport/, 2010. Accessed on May 22nd, 2010.

[32] Google, “Google AppEngine for Business Support.” http://code.google.com/intl/no-NO/appengine/business/support.html, May 2010. Accessed on May 22nd, 2010.

[33] S. Krishnan, “Geo Location Enables Developers To Choose Data Centers and Group Applications & Storage.” http://blogs.msdn.com/windowsazure/archive/2009/03/18/geo-location-enables-developers-to-choose-data-centers-and-group-applications-storage.aspx, 2009. Accessed on May 20th, 2010.

[34] Amazon.com, Inc, “Region and Availability Zone Concepts.” http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/index.html?concepts-regions-availability-zones.html, November 2009. Accessed on May 30th, 2010.

[35] Google, “Top ten advantages of Google’s cloud.” http://www.google.com/apps/intl/en/business/cloud.html, 2010. Accessed on May 14th, 2010.

[36] Google, “Comprehensive review of security and vulnerability protections for Google Apps.” http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en//a/help/intl/en/admins/pdf/ds_gsa_apps_whitepaper_0207.pdf, February 2007. Accessed on May 14th, 2010.

[37] K. Ward, “More Azure Hypervisor Details.” http://virtualizationreview.com/blogs/mental-ward/2008/11/more-azure-hypervisor-details.aspx, November 2008. Accessed on May 27th, 2010.

[38] H. Vo, “Design Principles Behind The Windows Azure Hypervisor.” http://blogs.msdn.com/b/windowsazure/archive/2009/01/29/design-principles-behind-the-windows-azure-hypervisor.aspx, January 2009. Accessed on May 27th, 2010.

[39] B. Hoard, “Hyper-V, We’ve got a Problem (Actually Three).” http://virtualizationreview.com/blogs/the-hoard-facts/2009/11/hyper-v-problem.aspx, November 2009. Accessed on May 27th, 2010.

[40] S. Lipner and M. Howard, “The Trustworthy Computing Security Development Lifecycle.” http://msdn.microsoft.com/en-us/library/ms995349.aspx, 2005. Accessed on May 27th, 2010.

[41] Windows Azure Team, “Introducing the Windows Azure Service Management API.” http://blogs.msdn.com/b/windowsazure/archive/2009/09/17/introducing-the-windows-azure-service-management-api.aspx, September 2009. Accessed on May 28th, 2010.

[42] EUSecWest, "Script runtimes are vulnerable just like everything else." http://eusecwest.com/justin-ferguson-interpreter-vm-attacks.html, May 2008. Accessed on March 15th, 2010.

[43] Google App Engine Team, "Google App Engine Blog: Happy Birthday." http://googleappengine.blogspot.com/2010/04/happy-birthday.html, April 2010. Accessed on April 12th, 2010.

[44] Stortinget, "Personal Data Act." http://datatilsynet.no/templates/Page____194.aspx. Accessed on February 16th, 2010.

[45] The Data Inspectorate, "Identity theft - An explorative study." http://www.datatilsynet.no/templates/Page____2603.aspx, 2009. Accessed on February 11th, 2010.

[46] C. Steel, *Core Security Patterns.* Upper Saddle River: Prentice Hall PTR, 2005. ISBN: 9780131463073.

[47] T. Limoncelli, *The Practice of System and Network Administration.* Boulder: Westview, 2007. ISBN: 0321492668.

[48] D. Swift, "A Practical Application of SIM/SEM/SIEM Automating Threat Identification." http://www.sans.org/reading_room/whitepapers/logging/practical-application-sim-sem-siem-automating-threat-identification_1781, February 2007. Accessed on May 19th, 2010.

[49] Microsoft Corporation, "Managing SSL Certificates." http://msdn.microsoft.com/en-us/library/ee758713.aspx, 2010. Last accessed on March 16th, 2009.

[50] Microsoft Corporation, "Security Guidelines and Limitations (SQL Azure Database)." http://msdn.microsoft.com/en-us/library/ff394108.aspx, 2010. Accessed on May 20th, 2010.

[51] Microsoft Corporation, "Crypto Services and Data Security in Windows Azure." http://msdn.microsoft.com/en-us/magazine/ee291586.aspx, 2010. Accessed on May 20th, 2010.

[52] BankID, "Tre av fem har tilgang til BankID." https://www.bankid.no/Presse-og-nyheter/Nyhetsarkiv/2010/Tre-av-fem-har-tilgang-til-BankID/, April 2010. Accessed on May 6th, 2010.

[53] Minister of Government Administration, Reform and Church Affairs, "Rammeverk for autentisering og uavviselighet i elektronisk kommunikasjon med og i offentlig sektor." http://www.regjeringen.no/en/dep/fad/dok/lover-og-regler/retningslinjer/2008/rammeverk-for-autentisering-og-uavviseli.html?id=505958, April 2008. Accessed on June 2nd, 2010.

[54] The Ministry of Modernisation, Norway, "Requirements specification for PKI for the public sector." http://www.regjeringen.no/upload/kilde/mod/rap/2004/0002/ddd/pdfv/250615-kravspekk-engelsk-versjon.pdf, January 2005. Accessed on May 6th, 2010.

[55] J. Ølnes, “Signaturkrav, autentisering.” http://ksikt-forum.no/filearchive/Jon%20Olnes.pdf, February 2008. Accessed on June 2nd, 2010.

[56] BankID, “BankID har høyeste sikkerhet.” https://www.bankid.no/Dette-er-BankID/BankID-far-du-av-banken-din/BankID-har-hoyeste-sikkerhet/. Accessed on June 2nd, 2010.

[57] Agency for Public Management and eGovernment (Difi), “Difi tildeler kontrakter om bruk av eID i ID-porten.” http://www.difi.no/artikkel/2010/04/difi-tildeler-kontrakter-om-bruk-av-eid-i-id-porten, April 2010. Accessed on May 6th, 2010.

[58] BankID, “Difi har avvist tilbud fra BankID Norge om tilknytning av BankID i felles infrastruktur for eID i offentlig sektor.” https://www.bankid.no/Presse-og-nyheter/Nyhetsarkiv/2010/Difi-har-avvist-tilbud-fra-BankID-Norge-om-tilknytning-av-BankID-i-felles-infrastrukt
April 2010. Accessed on May 6th, 2010.

[59] E. Zachariassen, “BankID skviset ut av eID.” http://www.tu.no/it/article242936.ece, April 2010. Accessed on June 2nd, 2010.

[60] I. Winkler, “Bank Theft Builds Case for Two-Factor Authentication.” http://www.internetevolution.com/author.asp?section_id=515&doc_id=181609&f_src=internetevolution_gnews, October 2009. Accessed on May 6th, 2010.

[61] Federal Financial Institutions Examination Council’s (FFIEC), “Authentication in an Internet Banking Environmen.” http://www.ffiec.gov/pdf/authentication_guidance.pdf, 2005. Accessed on May 6th, 2010.

[62] Norwegian Ministry of Trade and Industry, “LOV 2001-06-15 nr 81: Lov om elektronisk signatur (esignaturloven)..” http://www.lovdata.no/all/nl-20010615-081.html, June 2001. Accessed on June 2nd, 2010.

[63] A. K. Bentzen Ernes, “Mener BankID er trygg som e-signatur.” http://www.digi.no/498485/mener-bankid-er-trygg-som-e-signatur, November 2007. Accessed on June 2nd, 2010.

[64] Buypass, “Personal ID.” http://www.buypass.com/Home/Products+%26+services/Electronic+ID/Personal+ID. Accesed on June 2nd, 2010.

[65] Commfides, “e-ID – Elektronisk ID.” https://www.commfides.com/index.php?option=com_content&task=view&id=26&Itemid=48, January 2007. Accesed on June 2nd, 2010.

[66] BankID, “BankID er juridisk bindende.” https://www.bankid.no/Dette-er-BankID/BankID-far-du-av-banken-din/BankID-er-juridisk-bindende/. Accessed on June 2nd, 2010.

[67] C. Balding, “Cloudsecurity.org Interviews Guido van Rossum: Google App Engine, Python and Security.” http://cloudsecurity.org/blog/2008/07/01/cloudsecurityorg-interviews-guido-van-rossum-google-app-engine-python-and-security.html, 2009. Accessed on March 15th, 2010.

Appendix A

Feedback from The Data Inspectorate (Norwegian)

Datatilsynet

dSafe
Sem Sælands vei 5

7034 TRONDHEIM

Deres referanse	Vår referanse (bes oppgitt ved svar)	Dato
	09/01198-3 /SEV	20. november 2009

Avslutning av sak - Digital bankboks - dSafe

Det vises til virksomhetens e-post av 16. september 2009 og tilleggsopplysninger mottatt på e-post 1. oktober 2009.

Lovlighet
Produksjon og salg av produkter som benytter seg av eller er elektroniske systemer reguleres ikke av personopplysningsloven. Loven setter imidlertid begrensninger for hvem som lovlig kan ta i bruk slikt utstyr eller system. Det går videre et skille mellom bruk til privat formål og annen bruk. Dette skillet er videre drøftet under juridiske aspekter. Tilsynet tar videre opp spørsmålene rundt tekniske krav under punktet om tekniske aspekter.

Juridiske aspekter
Behandling av personopplysninger forutsetter oppfyllelse av grunnkrav etter personopplysningslovens § 11. Videre må den som iverksetter behandlingen påvise såkalt behandlingsgrunnlag etter personopplysningslovens § 8. Virksomheten har selv skrevet at dette skal være etter samtykke. Samtykke er definert i personopplysningslovens § 2 nr. 7: "En frivillig, uttrykkelig og informert erklæring fra den registrerte om at han eller hun godtar behandlingen av opplysningene om seg selv." I tillegg har virksomheten beskrevet sitt forhold til databehandler som skal oppbevare dataene på en sikker måte for virksomheten.

Virksomheten legger opp til cloud computing med en leverandør registrert utenfor Norge. I personopplysningslovens § 15 jf. § 13 er forholdet til en databehandler regulert, og Datatilsynet kan ikke se hvordan virksomheten skal overholde informasjonsplikten ovenfor kunden jf.§ 2 nr. 7 og kravet til informasjon i lovens § 19 når virksomheten er avhengig av en databehandler som vil kunne flytte personopplysninger fritt mellom land med ulik personvernlovgivning. Hva gjelder sistnevnte forhold vil det videre kunne oppstå konflikt med personopplysningslovens § 29 i den grad personopplysningene flyttes til land utenfor EØS-sonen.

Tekniske aspekter
Teknisk er ikke løsningen så komplisert i forhold til personopplysningsloven. Virksomheten må rette seg etter personopplysningslovens § 13 og personopplysningsforskriftens kapittel 2. Dokumentasjonen som virksomheten har oversendt ser med første øyekast ut til å kunne dekke disse.

Datatilsynet ønsker til slutt å presisere at dette er tilsynet generelle tolkning av personopplysningslovens bestemmelser i forhold til denne typen sikkerhetsprodukter. Det er den person eller virksomhet som tar løsningen i bruk som har ansvaret for behandlingen etter personopplysningsloven. Vedkommende er selv ansvarlig for å vurdere om kravene som er beskrevet over er oppfylt i det konkrete tilfellet.

Med hilsen

Leif T. Aanensen
avdelingsdirektør

Stein Erik Vetland

Stein Erik Vetland
overingeniør

Appendix B

Email Correspondence with the Data Inspectorate (Norwegian)

Til Datatilsynet:
Jeg skriver masteroppgave i informasjonssikkerhet og jobber for tiden med ISO27001-standarden. Jeg har lett etter informasjon angående hvordan Datatilsynet forholder seg til denne standarden når f.eks. et firma skal implementere og dokumentere sine sikkerhetspolicyer. Kan dere hjelpe meg med dette?

Svar:
Til Mats Andreassen

Datatilsynet forholder seg normalt til kravet om internkontroll i personopplysningslovens paragraf 14, jf. forskriftens kapittel tre. Men det er absolutt ingen bakdel å benytte ISO27001 for å standardisere sine policyer, men det er ikke et krav fra vår side ved en eventuell kontroll.

Med Vennlig Hilsen Frank U Eriksen Overingeniør, Datatilsynet

Til Datatilsynet:
Hei.

Takk for raskt svar på min forespørsel. Ja, jeg antok at ISO 27001 compliance ikke er krav fra deres side, men det jeg egentlig er ute etter er om dere har satt dere såpass godt inn i standarden til at dere kan finne på å /anbefale /det for firmaer som skal vise sin compliance med nevnte paragraf 14. Hvis dere faktisk anbefaler eller i det minste oppmuntrer til å benytte standarden vil det påvirke mine anbefalinger til det firmaet jeg behandler i min oppgave.

Svar: Vi oppmuntrer helt klart til å bruke det hvis virksomheten ser det som formålstjenlig med hensyn til størrelse, behov og eventuell sertifisering. Det vil helt klart hjelpe på å implementere internkontroll i en virksomhet med tanke på opplæring, avvikshåndtering og revisjon.

Frank

Til Datatilsynet:
Det virker veldig fornuftig. En må helt klart ta stilling til flere ting enn den nevnte lovparagrafen når en bestemmer seg for å gå for sertifisering, men de andre faktorene er opp til firmaet å vurdere. :)

Helt til slutt: kan jeg sitere deg i oppgaven min?

Svar:
Beklager. Det kan du :)