Fraunhofer
AISEC

FRAUNHOFER RESEARCH INSTITUTION AISEC

CLOUD COMPUTING SECURITY

PROTECTION GOALS.TAXONOMY.MARKET REVIEW.

DR. WERNER STREITBERGER, ANGELIKA RUPPEL

02/2010

Fraunhofer AISEC
Parkring 4
D-85748 Garching b. München

Tel.: +49 (0)89-322-9986-0
Fax: +49 (0)89-322-9986-299
http://www.aisec.fraunhofer.de

Published on 25 September 2009
State of study: September 2009

Contents

List of Figures

List of Tables

1 Executive Summary

The aim of this study on cloud computing security is to provide a broad framework for analyzing security problems in cloud computing systems. It is addressed at decisionmakers in enterprises in all branches of industry who already outsource IT services, use cloud services, or are considering deploying cloud services in the near future. It also targets anyone with an interest in the topic or wishing to gain an overview of the security risks arising from the use of cloud computing systems as well as the various systems currently available in the market, their costs, and their security concepts.

The study outcomes are presented briefly in the following, together with diverse security aspects that need to be taken into account by users of cloud services:

- The security and availability of cloud computing systems are two of the most important topics that must be considered by any cloud project. Almost every major vendor of cloud services has recorded a serious incident in one of these areas at some time in the past.
- Small and medium sized enterprises (SMEs) can improve their security by deploying cloud services: first, because they can procure security solutions as a service from specialist vendors and second, because they profit from the vendor's experience with implementing and operating secure services. On the other hand, the selection of a certified and trustworthy vendor, whose cloud services are delivered on the basis of a service level agreement that can be verified at any time, is a vital prerequisite.
- Large corporations should assess the cloud vendor's security functions individually and decide on a case-by-case basis whether the available security mechanisms are adequate for a particular application.
- The chief arguments in favor of cloud computing systems are the ability to leverage economies of scale and hence cut costs, the fact that capacities can be matched to current requirements, and the new opportunities for organizing existing processes.
- There are certain risks associated with the security and availability of cloud services, and there is a possibility of lock-in effects arising from the selection of a service. These can involve high costs, for instance if the service offered by a cloud vendor is switched and substantial changes are necessary to the existing system owing to a lack of standardization. The protection goals linked to security and availability resemble those of IT

security, namely confidentiality, integrity, authenticity, accountability, non-repudiability, availability, and the protection of privacy, and they must be defined when specifying the requirements.

- The protection goals of IT security can also be applied to cloud computing systems. However, they are too general to permit a precise analysis of these systems and their various forms, which means they have to be re-assessed and adopted for each individual cloud service. The main reason for this is the lack of standardized procedures for selecting and deploying security technologies in cloud computing systems.
- The structure of cloud computing systems comprises four layers – end user, software, platform, and infrastructure – and the players acting on these layers form a very complex IT security framework. This study describes all the key layers and players that must be examined, depending on the application and the selected cloud service.
- Certified procedure models as well as standardized interfaces and protocols based on cloud services are essential for cloud computing systems to increase the portability and interoperability of individual cloud service offerings. Standardization bodies, reference implementations, and development environments adapted to cloud computing systems must exist for this purpose.
- The cloud security taxonomy provides a clearly structured framework of the security areas that should be considered when using cloud services. Owing to the rapid development pace of both the technologies and the existing services, the application of the cloud taxonomy should be project based and the weighting accorded to individual security areas adapted to the specific requirements in each case.
- Modern cloud service portfolios clearly use a whole series of security technologies already, especially on the infrastructure layer. On the other hand, when it comes to architecture, administration, and compliance, cloud vendor support for security technologies is not yet adequate to achieve the stipulated protection goals. More detailed analyses are called for here to identify which current technologies are potentially suitable and determine whether new technologies need to be developed. There is a trend toward procuring certain security functions, such as parts of the identity or access management functionality, as a service from specialist vendors.
- On the administration side, service level agreements are an important instrument for specifying all the rights and obligations that exist between cloud users and cloud vendors. The standardized service level agreements offered at present, which are not normally freely negotiable by cloud users but can simply be either accepted or rejected, provide only minimal guarantees regarding cloud service quality. In particular, the security guarantees contained in these agreements are very rudimentary, and need to be extended in order to achieve the above-mentioned protection goals.

Systems to facilitate automatic monitoring and testing of the agreed service quality criteria are also essential.

- From the compliance perspective, there are no objections to the use of cloud services. However, the responsibility for the data concerned usually lies with the cloud user, who needs to define precise guidelines stating which information is allowed to be stored and processed in a cloud service and how, and simultaneously specifying the necessary security functions. From a legal viewpoint, too, the restrictions to which certain data is subject and the use of specific cloud services should be separately considered in each case.
- The market overview incorporated in this study gives a general rundown of selected cloud services together with their prices and functionalities. The taxonomy of secure cloud computing is then applied to these services and their security functions assessed. It is fair to say here that information about the implemented security functions is not adequately documented by cloud vendors. In many cases, security plays only a minor role in the presentation of their services, so that detailed information should be requested from the vendor upfront of choosing or using a specific cloud service. If appropriate, a proof of concept should be realized before the service is actually put to productive use.

2 Introduction

Cloud computing has developed into an increasingly important topic for many enterprises in the last few years, with the result that cloud services are already featured in a large number of end user applications [26]. The motivation for companies to consider cloud computing lies in the constantly evolving challenges that accompany the growing dynamics of the market and the ever fiercer competitive arena. It is consequently more vital than ever to continuously adapt and re-examine the know-how, the technology, and above all the internal resources that are employed.

The use of compute-intensive information technology (IT) is meanwhile an indispensable part of business operations, to enable business processes to be better targeted and new business solutions provisioned with greater flexibility and speed. The other side of this coin are the high costs for purchasing, operating, and maintaining the IT. These costs only rarely justify complete coverage of the maximum anticipated software and resource requirement, for example storage and computing capacity. In addition to improving efficiency and speed, enterprises therefore also have to realize cost savings and optimize the IT security of their infrastructure if they want to stay competitive.

Cloud computing can be the next step toward improving IT services and making better use of existing capacities. The concept that forms the basis for cloud computing describes various possible strategies to guarantee the dynamic deployment of IT resources, such as storage capacity or computing power as well as internal enterprise services or services across company boundaries. Cloud computing systems allow infrastructure resources and application services to be procured on demand as an IT service and thus outsourced to the cloud.

In the cloud computing paradigm, information is stored online on the same computers that are also used to run software applications [18]. These are made available to end users on request [27]. Although the information technology that provisions the data and applications is frequently operated by specialist vendors, the configuration is normally carried out by the end users in a web browser [23].

Since cloud computing systems are a paradigm that is continuously evolving, it is not possible to formulate a lasting definition of the term "cloud computing" at the present time[1]. A working definition drawn up by the National Institute of Standards and Technology (NIST), which is regularly updated and developed further, has been used for the purposes of this study [28]. The NIST defines cloud

[1] September 2009

computing as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e. g. networks, servers, storage, applications, or services) that can be rapidly provisioned and released with minimal management effort or (human) service provider interaction. This model promotes resource availability.

In addition to this, the NIST also defines the characteristics of, and the deployment and service models for, cloud computing systems. The five essential characteristics of cloud computing systems are outlined in this chapter while the deployment and service models are examined in detail in chapter 4. According to Mell and Grance [28], the five essential characteristics of which cloud computing systems are comprised are on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service (refer to Figure 2.1):

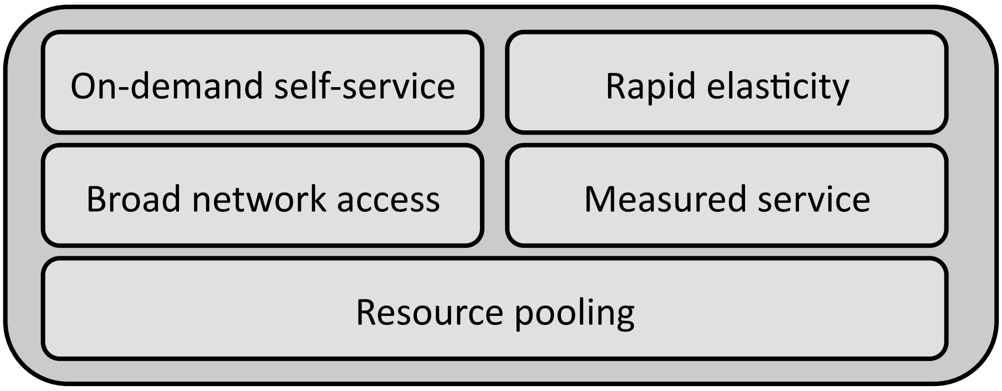

Figure 2.1: Essential characteristics of cloud computing systems

- On-demand self-service: A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.
- Broad network access: Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e. g. mobile phones, laptops, and PDAs).
- Resource pooling: The provider's computing resources are pooled to serve all consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e. g. country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, virtual machines, and service instances.
- Rapid elasticity: Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for rent often appear to be infinite and can be purchased in any quantity at any time.

- Measured service: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e. g. storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the vendor and consumer of the utilized service.

The five essential characteristics of cloud computing systems are applied in the following to grid and cluster computing systems, in order to distinguish them from cloud systems. All three systems are distributed and share similar characteristics that are summarized in Table 2.1. The similarities relate to resource pooling and broad network access – two criteria that are fulfilled by all systems. Network access to cluster and grid computing systems usually takes place within a corporate network, while the services of a cloud computing system can also be accessed through public networks like the Internet.

Table 2.1: Comparison of the three distributed system types: cluster computing, grid computing, and cloud computing

	Cluster	**Grid**	**Cloud**
On-demand self-service	No	No	Yes
Broad network access	Yes	Yes	Yes
Resource pooling	Yes	Yes	Yes
Rapid elasticity	No	No	Yes
Measured service	No	Yes	Yes

The differences between cloud computing systems on the one hand and grid and cluster computing systems on the other are attributable to the system dynamics. Resources in grid and cluster environments are generally prereserved, while cloud computing systems are demand driven, i.e. operation of these systems is geared to consumers' actual needs. Another difference concerns the "rapid elasticity" criterion, which forms an integral part of cloud computing systems but is not normally supported by cluster or grid systems. Service usage only tends to be accurately measured in grid and cloud computing systems, whereas the majority of cluster environments simply provision rudimentary metering functions.

Compared to other distributed systems such as grids or clusters, cloud computing solutions give enterprises significantly more flexibility. They can dispense with IT infrastructures of their own and only have to pay for the resources and services they actually use. These can be dynamically adapted to changed business requirements and processes with the help of virtualization technologies and service oriented, distributed software systems.

At the same time, the use of cloud computing systems also involves a number of security risks – most of them linked to the insufficient use of, and support for, security technologies. Yet-to-be-developed or immature technologies can likewise lead to security deficiencies in cloud computing systems [22]. As a result, the use of cloud computing systems is still restricted, and a detailed assessment of the potential security risks is essential because users expect secure cloud services to comply with the same high security standards as the systems used in the past. These risks can have a significant influence on the end user's business model – for instance, if confidential information is stolen.

According to a recent study by IDC[2], the advisory service provider, the security of cloud services is one of the most important reasons why cloud computing systems are not used in enterprises. Security is mentioned as a top priority criterion, alongside availability and costs, that must be satisfied before the complete breadth of cloud services can become a viable alternative to existing outsourcing concepts. Since only a few German companies have addressed this topic to date, it is likely that the significance attached to cloud service security will further increase in the future.

Table 2.2 shows the most important differences between classic outsourcing solutions and cloud computing systems as an alternative outsourcing concept. The term "IT outsourcing" describes the relocation of part or all of a company's IT to external vendors. Classic IT outsourcing can be anything from selective to total. Whereas with selective outsourcing, only specific IT functions are transferred to third-party vendors, total outsourcing covers the complete IT. The infrastructure and software can therefore be managed either by the customer or by the vendor, depending on the chosen configuration, although system administration is always the vendor's responsibility.

Table 2.2: Comparison of IT outsourcing and cloud computing

Characteristics	**Classic IT outsourcing**	**Cloud computing systems**
Technology location	Customer or vendor	Vendor
Business process adaptation	Vendor	Customer
Contract period	Medium to long-term	Short to long-term

Table 2.2 shows the three principal differences between classic IT outsourcing and cloud computing with reference to the "technology location", "business process adaptation" and "contract period" characteristics. With IT outsourcing, the infrastructure and software can be managed either by the customer or by the vendor. If a cloud computing system is opted for, both the infrastructure and the software are the vendor's responsibility. Another difference concerns the adaptation of applications to business processes. Whereas with classic IT outsourcing, the application is adapted to the business processes by the vendor, with cloud computing this is up to the customer because the cloud vendor is only responsible for provisioning and running the services. Lock-in effects are possible in connection with IT outsourcing, depending on the outsourced scope and the adaptation of the applications to the business processes, due to the

[2] More information about this study can be found in IDC's press release dated June 2, 2009: `http://www.idc.com/germany/press/presse_cloudcomp.jsp`

generally very long contract periods. Classic IT outsourcing is characterized by medium to long-term contracts while cloud computing systems typically feature short-term contracts.

The security aspects of a cloud computing system for different end user groups are briefly discussed in the following. We also examine whether it is possible to achieve the same IT security standard as with corporate security solutions. The most frequently cited advantages and risks of cloud systems are then described and the structure of the remaining study chapters presented.

2.1 Security aspects of cloud computing systems

The IT security of data, processes, and applications is one of the most important problem areas still linked to cloud services [26]. Until enterprises have access to mature security solutions that are adapted to, and support, the five essential characteristics of cloud computing systems, it will be very difficult for them to leverage the full potential of cloud services.

The use of cloud computing systems makes the security and availability risks increasingly opaque for cloud service users [37]. At the same time, the highly automated nature of cloud systems inevitably means a loss of control, so that cloud consumers have only limited influence on the geographical location of their data, for instance, or on the allocation of resources.

The increased use of cloud services gives rise to new weaknesses and threats on the IT security side that have to be considered when choosing the most suitable system. On the one hand, these new weaknesses can be attributed to attackers who assume the role of consumers in a cloud computing system in order to gain access to the data of other consumers. On the other, they stem from the complexity and dynamics of cloud systems, which are constantly changing due to outages or maintenance work. In addition, new methods for managing the risks must be evaluated, and the compliance of cloud computing systems with statutory requirements and guidelines needs to be verified. Only a small number of cloud vendors currently support the verification of processes according to predefined security guidelines [35].

This raises the question of whether cloud service usage is likely to entail a reduction in the level of IT security or whether cloud services can actually increase security. The different perspectives of the cloud service end user groups are crucial here. End users in small and medium sized enterprises often do not have the resources to draw up detailed security guidelines for their companies or to call on the necessary expertise to enforce them[3][4][5]. It could be argued that the

[3] http://www.ihotdesk.com/article/19055538/SME-security-could-benefit-from-cloud-computing-source-claims

[4] http://www.scmagazineus.com/sme-security-sme-mindset-must-change/article/136052/

[5] http://www.rationalsurvivability.com/blog/

existing security standard in this user group will be increased as a result of using cloud services, because implementing adequate security mechanisms is one of the cloud vendor's core tasks. It is assumed that state-of-the-art security technologies and the corresponding processes are realized by appropriately trained personnel on the cloud vendor's payroll.

However, this argument is offset by the fact that large companies in particular are very often motivated to use cloud services by the promise of cost savings; at the same time, cloud service vendors aspire to offer their services at the lowest possible price, which means forgoing certain security functions. If this scenario applies, the level of security could be impaired, potentially threatening the data, processes, and applications in the cloud. Enterprises must adapt their existing security systems, so that their security concepts also take account of cloud services.

New concepts and methods that are capable of identifying potential security risks and provide suitable technologies for containing threats need to be developed in both scenarios. Ideally, the option of using cloud services should be considered when the application is first designed, and security requirements complied with in all phases of the software development cycle. The costs generated by additional security mechanisms – be they services purchased from external vendors or instruments integrated directly in the applications – must be allowed for. At the same time, the lack of standardized security technologies and best practice approaches makes it more difficult to assess cloud service security. Security solutions for clouds will probably have to exhibit similar characteristics to the cloud systems themselves in terms of scalability, dynamics, fault tolerance, and openness in order to internalize the economies of scale that are facilitated by cloud services.

The fundamental challenges facing the security solutions for cloud services are attributable to the information asymmetries between cloud service vendors and end users. At the time of signing the contract, for instance, a cloud vendor knows more about the actual status of its system than the user of the cloud service. This gap is especially large in cloud computing systems because end users have very little information about, or influence on, how the services are provisioned and delivered by the vendor, whereas the latter has access to extremely detailed data. The security concepts for cloud services must strive to reduce these information asymmetries, for example by giving users access to monitoring and measurement data or by facilitating automatic testing of the vendor's security functions by a trustworthy third party who is qualified to assess the complexity of the cloud computing system.

The objective of this study is to provide an introduction to the structure of cloud computing systems and examine the above-mentioned question of whether these systems can provide increased security to cloud users. Based on this structure, we then show a possible breakdown of security areas into several categories encompassing all important aspects specific to cloud computing systems.

This novel framework can help companies identify the security risks more effectively and consider the strategic deployment of cloud services from the security perspective. This taxonomy is subsequently applied to selected, typical cloud services and their existing security functions analyzed.

2.2 Advantages and risks of cloud services

Enterprises that are toying with the idea of using cloud based services have to identify and understand the risks associated with them. This is essential in order to define detailed scenarios and implement risk management controls of the kind generally applied when handling confidential data or information that is subject to statutory regulations. Cloud computing systems entail the same risks as any other outsourced IT service. Questions such as data integrity, the recovery of data and processes, privacy protection, and special legal requirements are also particularly important in cloud computing systems and therefore need to be taken into account in the security analysis.

From the point of view of security and risks, cloud services oblige users to relinquish their traditional comprehensive control over data and processes by automating service provisioning, resulting in a gradual loss of transparency. In particular, the optimization of resources by the vendor can lead to unauthorized manipulation of customer data, as a result of which it is separately processed and archived at different locations.

These risks contrast with the opportunities created by the use of cloud computing systems, most of which are of a financial nature. By making use of cloud services, a company can improve its utilization of resources and the efficiency of its business processes while simultaneously increasing its IT flexibility. The following are frequently cited as key advantages of cloud computing systems:

- Reduced investment risks: The vendor bears the costs for purchasing the software or IT infrastructure components and thus also the investment risks, while the customer only pays for actual service usage or consumption.
- Improved performance and security: Specialized providers whose core business is operating the information technology generally have more resources at their disposal to guarantee the required performance and security. These additional resources can help improve the security standard.
- Scalable and flexible IT infrastructure: Cloud computing systems give enterprises an opportunity to add dynamic resources to their existing resources on demand and release them again when they are no longer needed. Possible project objectives no longer depend on the availability of sufficient server or storage capacity. Performance is often measured and optimized with the help of service level agreements.

- Lower costs of ownership: Cloud computing systems use methods already familiar from autonomic computing, such as self-healing. This increases their availability as well as their ability to be more self-managing. IT system administrators are no longer burdened by simple tasks and are free to concentrate on more complex activities.
- More efficient use of existing hardware and resources: Since cloud computing systems have a distributed architecture, they enable large amounts of unused IT infrastructure capacity to be leveraged throughout the company, so that purchases of new hardware are reduced to a minimum.

The incentive for cloud service vendors lies in the high scalability, leading to economies of scale through standardization. Although the fixed costs for setting up the infrastructure are higher, the variable costs for maintenance and support are substantially lower. From a technical viewpoint, cloud computing systems are supported by virtualization technologies and concepts that permit resources to be accessed in a shared pool.

These virtualization technologies divide physical IT resources into logical units that are made available to different end users and allow the resources to be simultaneously used. Instead of being exclusively assigned to a particular server or memory location, a cloud service is simply allocated a resource pool, such as a virtual machine, that is abstracted from the hardware. Additional free capacities are provided from a resource pool on demand. In other words, several logically distinct customers can be served via a shared infrastructure, so that the cloud service vendor’s infrastructure use is optimized.

2.3 Cloud computing scenarios

Cloud computing systems can be used extremely flexibly in different scenarios, as the work of the Cloud Computing Use Case Discussion Group shows [8]. The two scenarios of relevance for this study are described in detail later in this chapter together with their respective requirements. The scenarios are also presented in Figure 2.2. The first scenario illustrates the perspective of an end user accessing a cloud service, while the second scenario shows an enterprise that accesses and uses the resources of a cloud computing system.

Scenario 1: End user – cloud In this scenario, an end user is accessing data or applications in the cloud computing system. Common applications in this scenario include email solutions such as GoogleMail and social networking sites (e. g. Facebook or Twitter). These are accessed by end users through a web browser on almost any device, and each user is able to retrieve or manipulate their own data. The end users authenticate themselves to the cloud service with a user name and password; their data is stored and managed in a cloud computing system. Most importantly, the user has no idea how the underlying

Figure 2.2: The cloud scenarios: End user – cloud and enterprise – cloud

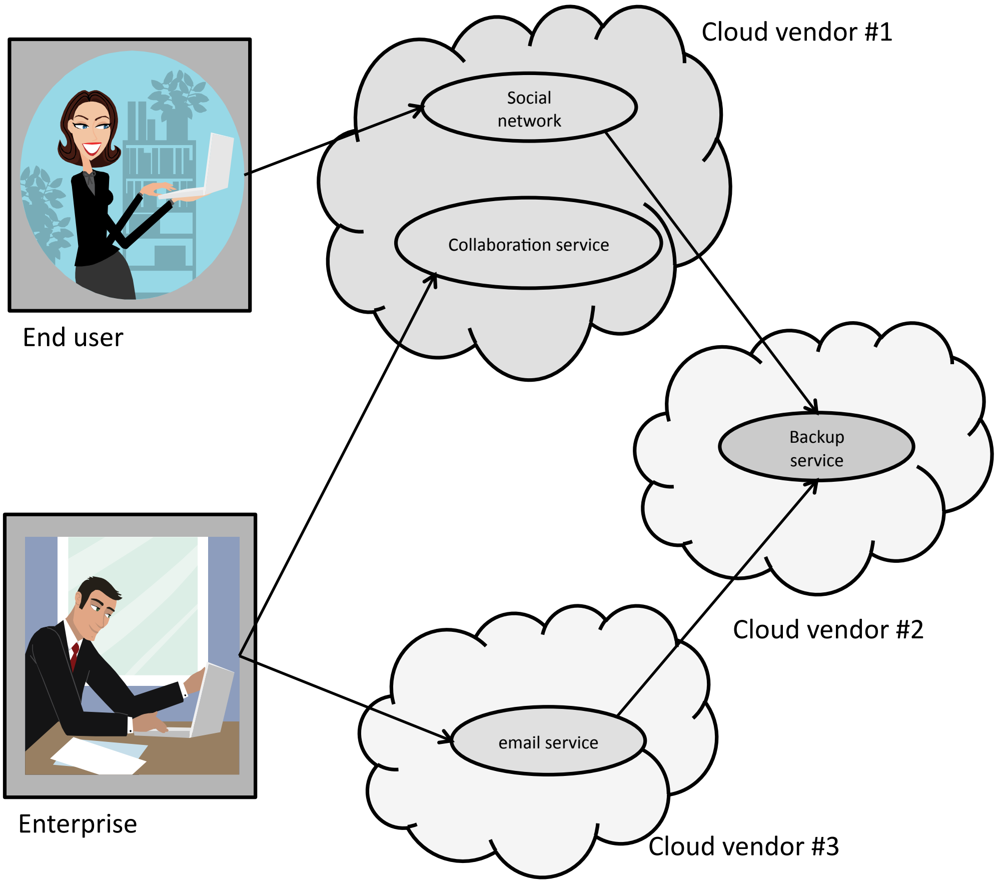

architecture works. If he or she can get to the Internet, they also want to get to their data – independently of geographical locations or technical restrictions, for instance.

The most important requirements in scenario 1 are as follows:

- Identity: The cloud service must authenticate the end user.
- An open client: Access to the cloud service should not require a particular platform or technology that restricts service usage.
- Security: A scenario between an end user and cloud computing services should include standard security techniques such as an encrypted connection to the cloud service, options for configuring privacy protection, and various other security aspects, which are examined in more detail in the framework of this study.
- Service level agreements (SLAs): Although service level agreements for end users will usually be much simpler than those for enterprises, cloud vendors must be clear about what guarantees of service they provide and what restrictions must be expected. End users should compare the SLAs for different cloud services prior to making their choice.

Scenario 2: Enterprise – cloud In this scenario, an enterprise is using cloud services for its internal processes. This is the most common use case in cloud computing [8]. Another variant of this scenario would be the use of cloud computing services by an enterprise with the aim not only of supporting its internal processes but also of providing these services to external players such as business associates or end users. The enterprise in scenario 2 could use cloud storage services for data backups, virtual machines in the cloud to bring additional processors online to handle peak loads, or applications in the cloud for certain enterprise functions (email, calendaring, CRM, collaboration etc.).

The most important requirements in scenario 2 are as follows:

- Identity and identity management: The cloud service must authenticate the end user. Since a user who belongs to a company is likely to have an identity within the enterprise, he or she should also use this identity to access the services of a cloud computing system. Other security requirements may need to be considered here, for instance to protect the user's privacy.
- An open client: Access to the cloud service should not require a particular platform or technology that restricts service usage.
- Location awareness: Cloud vendors take care of various management and administrative tasks on behalf of the enterprise, including the assignment of data and applications to the physical resources of the cloud computing system. This can lead to security-critical requirements, for example if data crosses international borders. For this reason, the enterprise should always be in a position to track the geographical location of the datacenter in which its data and applications are archived.
- Metering and monitoring: All cloud services must be metered and monitored to enable consumption to be billed and breaches of contract or security problems detected.
- Security: As already mentioned, security represents a major challenge for any cloud computing system. Security requirements must take account of the five essential characteristics of these systems. This study focuses on the security aspects of cloud computing systems, which are described in detail in a later section.
- Interoperability and portability: Applications, data, and virtual machines should be portable between the cloud vendor's various cloud computing systems. This presupposes a uniform set of interfaces, standardized as far as possible, for accessing cloud services such as storage, middleware, or platform services. The aim is to avoid lock-in effects and connect the cloud services of different vendors together.

- Distribution: The distribution of the applications and data in a cloud computing system is closely linked to their interoperability and portability. Distribution requirements can also be externally specified in the form of compliance guidelines.
- Service level agreements (SLAs): In addition to the basic SLAs required by end users, enterprises will need a standard procedure for benchmarking SLA performance (refer to "Metering and monitoring"). There must be an unambiguous way of defining what a cloud provider will deliver in an SLA, and there must be an unambiguous way of measuring what was actually delivered.
- Lifecycle management: Enterprises that use cloud services must also be able to manage the lifecycle of their applications, data, or identities. Suitable processes and security mechanisms to support this requirement and implement it in a reconstructable manner must exist for this purpose.
- Governance: Public cloud providers make it very easy to open an account and begin using cloud services. This ease of use creates the risk that individuals in an enterprise will use cloud services on their own initiative, for instance to transfer sensitive data to a cloud computing system. Governance requirements can also affect the security aspects of cloud computing systems and should therefore be taken into account in the security concept.
- Industry standards and protocols: If existing systems are operated using cloud computing resources, the requirements of existing industry standards and protocols must be considered. Industry-specific requirements are not discussed any further in the framework of this study because the topic is far too complex.

These two scenarios and their requirements form the basis for a more detailed examination of the security aspects of cloud computing systems in the next few chapters. The focus is on the end user view, as described above. In the next section, we introduce you to the top ten dos and don’ts of cloud computing security that should be remembered by any enterprise considering using cloud services.

2.4 Cloud computing security: The top ten dos and don’ts

The top ten dos and don’ts of cloud computing security combine the most important activities and processes linked to the use of cloud services. They should always be remembered in order to leverage the full potential of a cloud computing system while simultaneously reducing the security risks to a manageable level.

Die 10 Do’ and Dont’s sind:

1. Use a holistic security concept: Cloud computing systems are complex, distributed systems comprised of a large number of components and services on different layers. Cloud services should therefore be evaluated according to the security aspects of the Fraunhofer AISEC taxonomy for cloud computing systems, as described in this study, in order to obtain a holistic view of IT security in the cloud system. It may be necessary to analyze the infrastructure, application, platform, administration, and compliance aspects in more detail, depending on the specific project.

2. Integrate the services in an existing security concept: Cloud services should be integrated in an existing security concept, and suitable measures implemented to apply and enforce this concept. Existing systems must be adapted to cloud systems for this purpose, for instance to ensure that central administration of the IT systems continues to be supported.

3. Build a relationship of trust between the cloud consumer and the cloud vendor: Owing to the highly automated nature of cloud computing systems, human interaction between the cloud consumer and the cloud vendor is no longer inevitable. For this reason, cloud consumers should arrange a meeting with the cloud vendor prior to using a cloud service, in order to put themselves in the picture directly about the provider's datacenters, employees, and processes. The persons who can be contacted in the event of problems should also be agreed, as a way to improve the relationship of mutual trust between the consumer and the vendor upfront of cloud service usage.

4. Protect the network infrastructure: Cloud computing services are always procured via a network – often the Internet – which means that the security and reliability of the network infrastructure require particular protection. Standard methods such as firewalls, encryption, and virtual private networks (VPNs), or redundant network connections should be taken into account here. An encrypted connection should be used for all communications with the cloud vendor.

5. Use innovative security solutions for cloud computing systems: The security solutions for cloud computing systems need not necessarily be purchased as a software product or developed in-house; they can often also be rented from external vendors who offer solutions specifically tailored to the characteristics of cloud computing systems. The aim should be to achieve end-to-end security that allows all end user accesses and actions on the cloud computing system – both by the cloud consumer and by the cloud vendor – to be reconstructed. This is especially important if other players who are unknown to the enterprise or the end user are involved in the provisioning and use of a cloud service.

6. Use basic services: Cloud services generally include a set of basic services linked to security, distribution, provisioning, or integration that can be procured over the platforms of the respective vendors, and these services

should also be used. A high security standard can normally be realized quickly and with minimal effort using these basic services. It is important to check that a security certificate exists for the services as proof that they have been security tested by a third party.

7. Pay attention to lock-in effects: The use of industry standards and open protocols simplifies the interoperability and portability of data and applications in a cloud computing system. However, there are currently no standards for cloud computing systems, and it is still not clear which technologies will become established in the medium to long term. Lock-in effects, linked to high switching costs, are therefore possible if a consumer switches to another cloud vendor.
8. Request security certificates: Asking for security certificates can be a pointer to the security of a cloud computing system. It is important to examine individually which processes have been security tested by an external company and how these processes are implemented by the cloud vendor. A security certificate issued by a third party should be requested for cloud services that provide the enterprise with critical functions. The certificate should confirm that a secure software development cycle has been used, for instance, or a penetration test carried out.
9. Don't forgo security concepts for purely financial reasons: It is not advisable to dispense with security concepts, as stipulated in the corporate security guidelines, for example, for purely financial reasons.
10. Use service level agreements: Service level agreements, in which all rights and obligations of the stakeholders must be defined, are central to the use of cloud services. The SLAs offered by the majority of cloud vendors as standard should be subjected to a critical scrutiny, and individual SLAs negotiated with the cloud vendor if necessary. In addition, suitable systems should be installed for monitoring the SLAs (automatically if possible), and the results measured by these systems submitted for regular compliance checks.

2.5 Structure of the study

Now that this chapter has provided an introduction to cloud computing systems and described the problems involved, chapter 3 will set out and define the protection goals. The scope that needs to be supported by a cloud service's security functions is determined by specifying the protection goals in as much detail as possible.

Chapter 4 describes the structure of a cloud computing system with the help of a layer model, as well as the most important players and various other parameters that influence cloud computing security. Building on the outcomes of chapter 4, chapter 5 derives a taxonomy encompassing all the most important security areas of cloud services – from architecture and infrastructure through cloud service management to cloud computing compliance.

Chapter 6 begins by outlining a few popular cloud services, then evaluates their security functions with reference to the cloud taxonomy. The chapter concludes with an assessment of the security functions. Finally, chapter 7 summarizes the study findings and provides an outlook to future developments in cloud service security.

Figure 2.3 shows the protection goals, technical abstraction level, threats, and protection strategies that will be discussed in the next few chapters. The technical abstraction level corresponds to the layer model in chapter 4, while the threats refer to the security areas explained in chapter 5.

Figure 2.3: Overview of cloud computing security

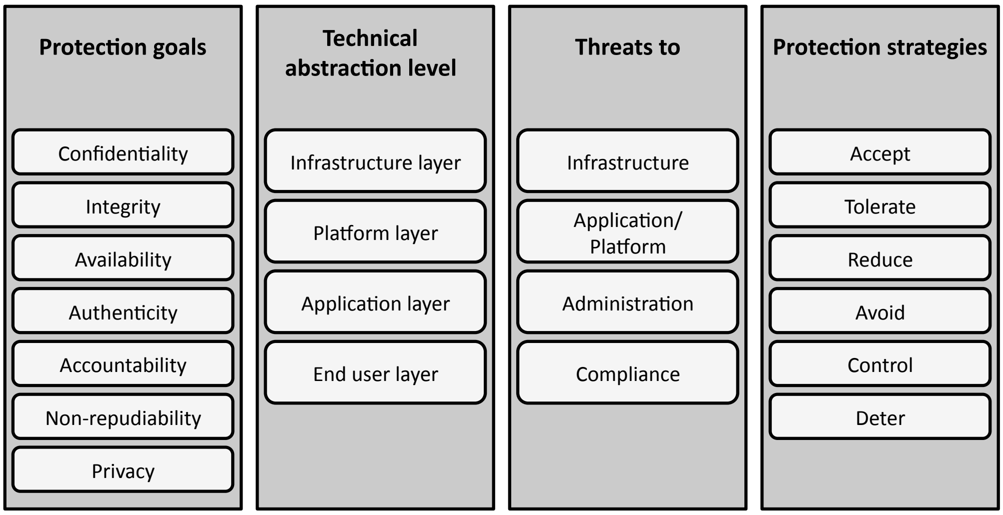

Protection goals	Technical abstraction level	Threats to	Protection strategies
Confidentiality	Infrastructure layer	Infrastructure	Accept
Integrity	Platform layer	Application/ Platform	Tolerate
Availability	Application layer	Administration	Reduce
Authenticity	End user layer	Compliance	Avoid
Accountability			Control
Non-repudiability			Deter
Privacy			

3 Protection goals

The protection goals form the basis for the security requirements that must be fulfilled by IT systems in general and cloud computing systems in particular. These goals are usually fixed for a specific scenario in the framework of the requirement definition and are part of the nonfunctional requirements to be met by the cloud service vendor as well as by the cloud service itself.

The six most important protection goals – confidentiality, integrity, availability, authenticity, accountability, and pseudonymity – are introduced in the next few sections, then explained in more detail with reference to selected cloud computing scenarios. Depending on the scenario concerned, individual protection goals can be accorded a higher weighting, for instance if confidential data needs to be stored, or they may play only a minor role, say, for running test systems in the cloud. The concept of multilateral security, which takes account of the protection interests of all stakeholders and the settlement of protection conflicts arising from these interests, for example in connection with the use of a cloud service, can be applied here.

3.1 Confidentiality

The confidentiality of a system is guaranteed providing it prevents unauthorized gathering of information [17]. In data secure systems, the "confidentiality" characteristic requires authorizations and checks to be defined, to ensure that information cannot get into the possession of subjects who do not have the appropriate rights. This comprises both access to stored data authorized by users and data that is transferred via a network. It must be possible to assign and withdraw the rights that are necessary to process this data, and checks must be implemented to enforce compliance. Cryptographic techniques and access controls based on strong authentication are normally used to protect confidentiality.

The data in a cloud computing system is very often in motion owing to the system's dynamic and open nature. A cloud resource vendor must be able to store this data on a server of its own choosing – and possibly also allowed to copy or duplicate it – in order to optimize its infrastructure capacity and ensure the necessary performance. These processes are usually outside the customer's sphere of influence and can lead to confidentiality problems, for instance if the data crosses territorial borders or is stored on a less secure system. In addition, the algorithms and data structures employed mean the vendor cannot always guarantee the data's availability on a storage medium in encrypted form. Moreover, the majority of cloud vendors fail to provide any assurances in their terms and conditions of business about where data is stored or the measures taken to protect its confidentiality [20]. In many cases, it is actually up to the customer to implement suitable security techniques. Data at rest should always be encrypted before it is archived on a storage medium or in a database. This includes internal enterprise information, data belonging to public authorities, personalized data, and other confidential information or data subject to statutory controls, such as credit card numbers.

A typical cloud scenario tends to involve not just one consumer and one vendor in a bilateral business relationship but a series of other vendors playing a variety of roles, for example as intermediaries or consumers of other cloud services. Whereas in the first instance – a bilateral business relationship – confidentiality can be assured using existing methods for secure data transmission like SSL/TLS, the second case necessitates broad support for technologies that guarantee confidentiality between a group of stakeholders. In addition to vendor guidelines describing the use and verification of confidential data, this also covers support technologies for managing the data encryption and decryption algorithms.

The management of the rights that are required in a cloud system to achieve the protection goal of confidentiality likewise create new challenges. Here, too, the chief problem is developing efficient methods for administering such a large number of players. In traditional enterprise architectures, data is generally protected by building a security zone in the form of a firewall that prevents access by potential attackers. A clear separation of rights inside the firewall from rights outside of it is vital. The data in a cloud is distributed across several systems that can have different geographical locations and be operated by different vendors. This scenario calls for new methods of accessing data and systems in order to comply with the protection goal of confidentiality.

3.2 Integrity

A system guarantees data integrity if it is impossible for subjects to manipulate the protected data unnoticed or in an unauthorized way [17]. Data, messages, and information are considered to have integrity if they are trustworthy and cannot be tampered with. A cloud computing system assures the integrity of the protected data if this information cannot be modified by third parties. If integrity is specified as a protection goal for cloud services, not only the cloud surface itself that is accessed by the end user must achieve this goal but also all other components with a stake in the cloud. In a complex, distributed system such as cloud computing, this can be a highly complicated task and is the responsibility of the cloud service vendor.

Data that is stored on a virtual hard drive, for instance, must be protected against unauthorized manipulation either by other participating systems used to process the information or by external attackers. Errors in the configuration of a cloud vendor's systems can also cause integrity to be violated, so that the data in a cloud computing system should always be provided with a cryptographic checksum. The original checksum can be stored on a trustworthy third-party computer for comparison purposes. The checksum should also be verified each time the data in the cloud computing system is accessed.

Software, configuration, and message integrity are likewise essential in a cloud system alongside data integrity. Software integrity ensures that the software used to run a cloud computing system is intact when it is delivered by the software manufacturer, in other words that it has no "back doors", for example, and has not been tampered with in any other way. Configuration integrity prevents the configuration of a cloud resource or a cloud service from being changed by unauthorized persons. This is particularly vital in cloud systems because cloud environments are normally automatically launched and managed by means of configuration scripts.

Since cloud computing systems are a kind of distributed system, message integrity is another key requirement that must be satisfied both within the cloud and between different clouds and the end user systems. In particular, the administrative and control information of cloud systems needs to be specially protected because these messages are often transported via public networks.

If several cloud services are involved when an end user deploys a complex service, integrity violations can occur if one or more cloud services cannot be run. The problem is further exacerbated if some of the cloud services concerned support transactions while others do not. Transactions in distributed environments like cloud computing systems serve to keep the actions of several stakeholders consistent. Protocols based on the all-or-nothing principle are normally used for this purpose, in other words any changes or calculations made are only stored persistently if the services are run successfully. Those cloud services that do not support transactions must be restored to their status prior to partial execution in the event that they are unsuccessful, in order to ensure the integrity of the data.

Cloud services that use XML based interfaces – such as web services based on SOAP or REST – are frequently procured with the Hyper Text Transfer Protocol (HTTP). At the protocol level, HTTP supports neither guaranteed delivery nor transactions, so that functions to achieve the protection goal of integrity must be implemented on the application layer.

3.3 Availability

DIN 40042 defines availability as the probability that a system will operate satisfactorily at any point in time. A cloud computing system should allow its users to access the required resources in the agreed way at all times. Its availability must not be restricted by unauthorized actions or targeted attacks by external players.

This protection goal presents cloud computing systems with a major challenge, because they are generally reached via a public network and hence exposed to the typical risks of all such networks, such as distributed denial-of-service attacks. In particular, errors in the system configuration or an excessive number of cloud service requests placing a heavy burden on the cloud vendor's infrastructure and impairing not just a single service but the entire cloud computing system have restricted the availability of cloud services on numerous occasions in the past.

The use of cloud computing systems causes the emphasis of strategies to assure high availability to be shifted from measures at the hardware level (e. g. redundant power supplies) to software measures. The reason for this is that mainly standard hardware components are employed and interconnected in large farms. While this has the effect of reducing the infrastructure vendor's capital costs, it also increases the probability of a hardware defect, which must be compensated by means of suitable software mechanisms.

Technical solutions can activate checkpoint and recovery mechanisms, for instance, to restore the status after an outage or support different redundancy based techniques. External attacks on the availability of the cloud computing system, such as the distributed denial-of-service attacks mentioned above, are generally restricted by limiting the resources provided to a single user, or else their impact is minimized by changing the network configuration. Both the cloud service vendor and the cloud user must be aware of these risks and implement suitable strategies for combating them while simultaneously guaranteeing maximum availability.

3.4 Authenticity

The authenticity of a subject or object is defined as its genuineness and credibility; these can be verified on the basis of its unique identity and characteristic features [17]. Information is authentic if it can be reliably assigned to the sender, and if it can be proved that this information has no longer been changed since it was created and distributed. A secure technique for identifying the communication partners and mechanisms for ensuring authenticity are essential here. These mechanisms must be capable of confirming or disproving the authenticity of the protected information. None of the system participants can create or distribute messages and data on behalf of another subject.

When an enterprise first begins to use cloud services, ensuring the authenticity of end users is a key requirement. Various identity management problems of a general nature have to be tackled, such as the administration of credentials, sufficiently strong authentication mechanisms, and the management of trust relationships between cloud services as well as across different cloud computing systems.

Digital signatures, security tokens, or passwords, which enable the signatory of a message or the creator of a signature to be identified, are normally used to verify authenticity in a cloud computing system. Federated identity management concepts based on attributes, which are usually procured in a distributed way from different identity vendors, are also possible. The aim is to guarantee the authenticity of all communication partners in the system.

The authentication procedure between a cloud user and a cloud service can be built around the exchange of authentication data, which can take place separately from, and independently of, the transfer of the application data. In a cloud computing system, not only the cloud user needs to be authenticated with the cloud service but also the cloud service with the cloud user. This prevents possible man-in-the-middle attacks, and stops data being transferred and processed by malicious cloud services.

3.5 Accountability

The protection goal of accountability requires actions to be clearly assignable to an actor in the system and ensures that the authorship of an event or action in the system cannot be rejected. All actions in a cloud computing system should be attributable to a player, even if this can result in the violation of a contract. For this reason, accountability always includes the identity of the action's author and a time stamp as well. It is extremely important for the binding legal force of electronic business transactions such as the use of a cloud service. When a cloud service is accessed, the protection goal of accountability ensures that all actions have been verifiably executed – in particular, vis-à-vis third parties – by a specific actor in the cloud computing system and that, as a result, they can be taken as a basis for billing resource usage, for instance.

Service level agreements that specify certain performance guarantees are a prerequisite for achieving the protection goal of accountability. These guarantees must be monitored by suitable systems and any variances documented. All other actions by the players in a cloud computing system must be additionally logged to allow them to be unambiguously assigned.

The accountability of a cloud service can be ensured, among other things, by means of qualified signatures, encryption, or mechanisms to protect data integrity. The non-repudiability process can usually be divided into four phases, which are specified in detail by a non-repudiability record: proof construction, proof transfer and storage, proof verification, and conflict resolution. In a cloud computing system, proof might be constructed using digital signatures that can be validated by a third party, for instance.

Let us assume that a resource vendor changes the guidelines for assigning the virtual machine's resources. This impacts the performance of the application a service vendor provisions to a service consumer. The change in capabilities can lead to a violation of the service level agreement that has been concluded between the service vendor and the service consumer. The protection goal of accountability must establish that the SLA violation is attributable to a change in the resource assignment.

3.6 Pseudonymity and privacy protection

The protection goal of pseudonymity serves to protect the privacy of persons. An IT system that protects the privacy of its users should only collect and store as much data about them as is actually required to provision the service, and it should only make this information visible to authorized persons. The technical and organizational measures employed for this purpose should ensure that no profiles can be created of use patterns. The anonymous use of services is privacy in the strictest sense of the word.

The protection goal of anonymity can only be partially achieved for cloud users because detailed profiles of their actions have to be created in order to bill the resources used. For this reason, cloud computing systems should implement pseudonyms that allow an actor (e. g. a consumer or a vendor) to reveal the identity hidden behind the profiles for billing purposes. In combination with the protection goal of accountability, this permits key privacy elements such as transparency, assurances, or compliance with guidelines to be monitored [30]. Machine readable guidelines to protect privacy are required for this purpose; their ability to achieve the protection goal must be measurable on the application layer, preferably independently of the application's implementation.

Only anonymous data must be made available to unauthorized users regardless of the cloud architecture's individual layers. When choosing a suitable vendor or services, attention should be paid to the processes employed to accomplish this objective. If different vendors are used, it is important to ensure that privacy is also protected from one vendor to another.

The end users of a cloud service, such as a social network on cloud resources, are often unaware that their data is stored on a cloud from which subsequent use is possible, for instance, with the result that their privacy is violated. In this scenario, it is essential for the users of a service to be given control over their data, so that this service can be accessed transparently.

4 Structure of cloud computing systems

This chapter provides an introduction to the structure of cloud computing systems and describes the most important layers, actors, and deployment and service models mentioned in the NIST's extended definition of cloud computing [28]. The aim here is to establish a common understanding of cloud computing concepts and to outline the threats created by cloud computing systems for the protection goals summarized in chapter 3. A taxonomy of cloud security risks is then derived in chapter 5 from the structure and threats presented here.

Figure 4.1: Characteristics and deployment / service models of cloud computing systems

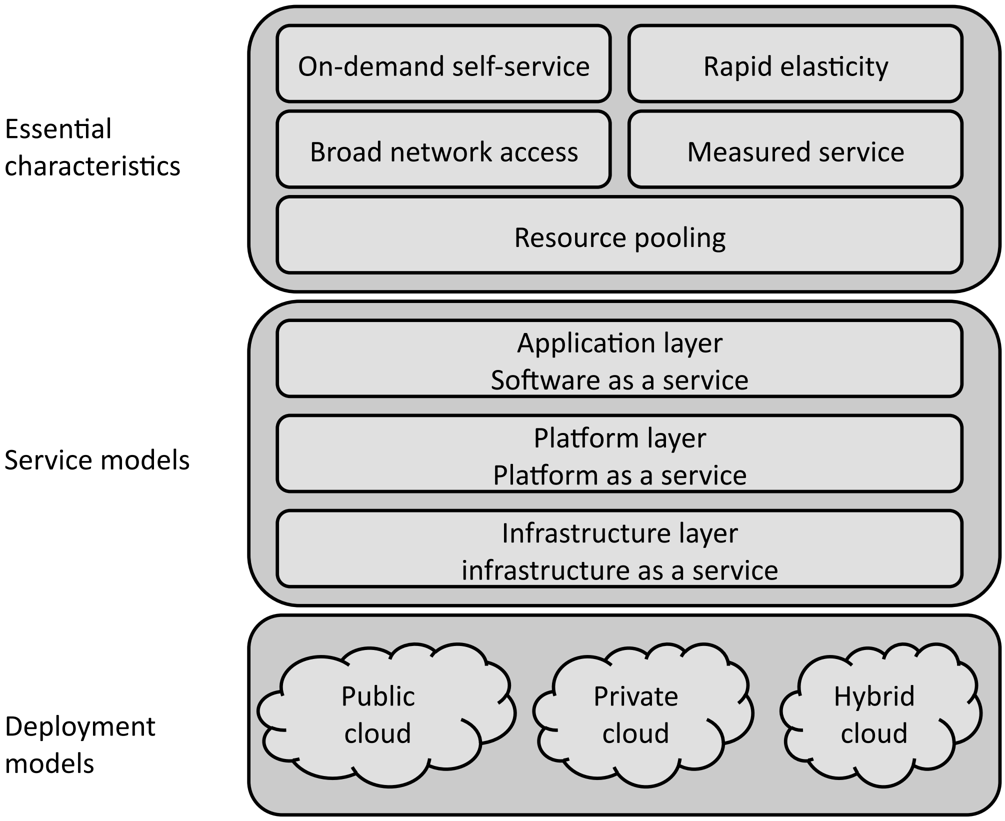

Figure 4.1 refers to the characteristics discussed in chapter 2 and extends them with the most frequently used deployment and service models for cloud computing systems, on which the structure of these systems is based. The various layers and variants of the deployment and service models are described in detail in the next few sections.

Section 4.1 explains the layer model for cloud computing systems, which is central to the service models and based on the services offered by these systems. There are three different layers for infrastructure, platforms, and applications.

In addition to the standard service model division into three layers, this study also introduces the concept of the end user layer, encompassing all players and systems that are relevant for consumers. Reference is made to the cloud computing scenarios formulated by the Cloud Computing Use Case Discussion Group [8]. The most important actors on each of these layers are presented, threats to security analyzed, and typical services described.

Section 4.2 takes a look at the model's end user layer. Depending on the type of cloud computing system, the end users on this layer use one or more of the underlying services on the application layer in section 4.3, the platform layer in section 4.4, or the infrastructure layer in section 4.5. Finally, various cloud service deployment models are discussed in section 4.6, which also considers the use of cloud services via publicly accessible or private cloud computing systems as well as hybrid models.

4.1 Layer model

The layer model used in this study is based on service offerings that are commercially available using cloud resources. These services can be divided into three categories, namely infrastructure, platform, and application. The three standard layers are augmented by an end user layer that encapsulates the end user perspective on cloud services [8]. The model is shown in Figure 4.2.

The end user view can incorporate different types of cloud services, which are accessible on the sublayers, in this model. If cloud users access services on the infrastructure layer, for instance, they can run their own applications on the resources of a cloud infrastructure and remain responsible for the support, maintenance, and security of these applications themselves. If they access a service on the application layer, these tasks are normally taken care of by the cloud service vendor. Each cloud user can thus deploy the cloud services flexibly and explicitly, depending on their individual requirements. At the same time, they are forced to relinquish control over the geographical location of their data and applications.

Figure 4.2: Layer model for cloud computing systems

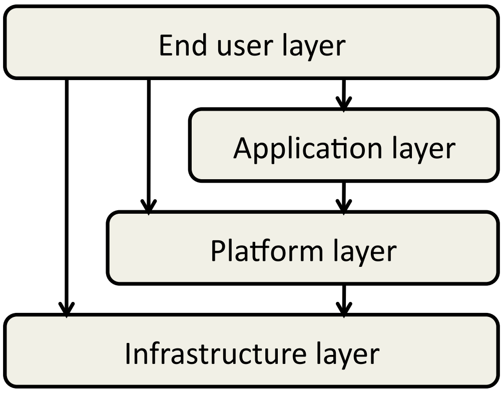

The layers shown in Figure 4.2 are described more extensively in the next few sections; their most important players are discussed and the threats that can arise to the protection goals set out in chapter 3 are analyzed.

4.2 End user layer

The end user layer of the cloud model comprises all systems, components, and devices that allow an end user to access cloud services on the sublayers. This access can be effected either automatically by the user's existing systems or manually by selecting the service via a portal. The adaptation of the legacy security functions for use with cloud services represents a particular challenge on the end user layer, and is the object of current research [9]. The findings published so far fail to state conclusively whether it is sufficient simply to extend the old solutions or whether new security technologies need to be developed for cloud computing systems, especially those with high security requirements.

One typical example that is very widespread in enterprise networks concerns identity and rights management. Many solutions in this area do not yet support the integration of external cloud services, so that either the existing systems must be extended or new solutions implemented to replace them.

The players on the end user layer are either software agents, which are integrated into the end user's systems and applications and generally act on the basis of predefined guidelines, or human users, who act as consumers of cloud services. The cloud service consumers, in turn, can be either corporate employees or private users; it can be assumed that the number of private cloud service users is significantly higher than the figure for business users, as the use statistics for cloud based services offered by Google (e. g. GoogleMail) and social networks (e. g. Facebook) show. In a cloud computing system, the end users are subscribers or tenants who purchase services from a service vendor within the system and are billed according to actual consumption. As a result, cloud users incur no, or only minimal, capital costs and are able to rent elastic, almost unlimited server and storage capacity. These capacities can be accessed from a whole series of mobile and stationary devices because they are usually provided as web services.

From the end user's point of view, there are several advantages to using cloud services, the majority of them linked to system support and maintenance, which in a cloud computing system tend to be the responsibility of the service vendor. It is important to assess precisely which routine administrative tasks are to be performed by the cloud service vendor and which activities must be taken care of by the end user. By transferring routine tasks to the cloud service vendor, the end user can deploy the freed-up capacities to develop innovative applications, for instance, or to concentrate on the enterprise's core competencies.

At the same time, it is up to each end user to weigh up the pros and cons of cloud service usage. It is imperative that they build up a certain expertise in risk assessment – especially security and reliability risks – and that they also consider monetary aspects in their decisions. The perspective of the end users in an enterprise reveals a shift in the requirements profile of IT staff like administrators away from routine tasks – such as system maintenance and support – and toward risk and cost evaluations linked to the use of cloud services.

Private users can profit from cloud services in a similar way in that they no longer have to bother with application maintenance and support, yet can still take advantage of the latest application version as a web browser based service. Further benefits for private users derive from the simple enabling and shared use of data, which can be made accessible to a defined circle of friends, for instance.

4.3 Application layer: Software as a service (SaaS)

The application layer is the layer in the cloud model that is visible to a cloud's end users and whose services are normally deployed by end users. It is mainly accessed through a web portal and service oriented architectures based on web service technologies. Credit card or bank account details must be provided to enable the fees for the use of the services to be billed.

The services on the application layer can be seen as an extension of the ASP (application service provider) model, in which an application is run, maintained, and supported by a service vendor. The main differences between the services on the application layer and the classic ASP model are the encapsulation of the application as a service, the dynamic procurement, and billing by units of consumption ("pay as you go"). However, both models pursue the goal of focusing on core competencies by outsourcing applications.

Since cloud applications are provisioned on the vendor's infrastructure rather than on an end user's server, they generate a whole set of advantages for both the service consumer and the service vendor on the application layer. Service consumers are generally actors who access a cloud service and are bound by the mandated rules of the organization (e. g. an enterprise or a department) to which they belong. These rules can result in minimum requirements being specified for the cloud vendor with regard to functional or nonfunctional criteria, such as service quality or the security functions the vendor must implement for a particular cloud service.

In exceptional cases, service consumers can also represent a threat to the cloud service, for instance if they send an excessive number of service requests to the system, so that the latter is no longer able to cope with the processing volume, or if they exploit a service's weaknesses in order to infiltrate it with foreign code.

Service vendors on the application layer offer an application service that is provisioned using the platform layer or the infrastructure layer. Service vendors with no infrastructure of their own rent it from a resource vendor.

Central provisioning of services by a service vendor for a large number of service consumers has the advantage – particularly on the application layer – of allowing more efficient maintenance as well as short innovation cycles, because the services are provided on a known, extremely homogeneous platform and no longer need to be tested on a multitude of different system configurations. Cloud services can thus always be made available to the service consumers in the latest version and integrated into existing processes without difficulty with the help of a service oriented architecture.

In spite of the numerous benefits of cloud services on the application layer, the security functionality and availability remain the most notable problem areas that restrict secure service deployment [37]. Secure service access must be guaranteed not only between the end user layer and the application layer but also between services on the application layer and other services on the sublayers of this layer. If a set of services from other service or resource vendors are integrated, the outage of a partial service can have a significant influence on service provisioning on the application layer. In the worst case, it may no longer be possible for the service to be accessed by the end user. Cloud service end users must be aware of the risks involved for this reason; they must also seek information from the vendor of an application service, then discuss possible scenarios and their consequences with the chosen cloud service vendor.

Cloud services in various application categories are already available today under the heading "software as a service". The most important application categories are as follows [38][25]:

- Scalable websites: This category comprises applications that are accessed over the Internet and have a large number of users. Load balancers, distributed databases, and scalable portals are components frequently implemented in these applications. Web based Office applications, social networks, and web based desktop environments are typical examples.
- Data management and distribution: Services for storing data and distributing it worldwide through content delivery networks (CDN) are commonly used in conjunction with scalable websites. Online backup services are another kind of data management application.
- Software development and testing: Cloud services in the software development domain offer multiple opportunities for collaboration; for example, it is possible to rent a shared directory for data interchange, source code management, and other services as a service. Cloud resources in the software testing domain can be used to create a global, distributed testing environment and carry out scalability tests.

- Interactive, mobile applications: This category provides a cloud service's application logic to the user of a mobile device. A cloud service can provision the data and the application very close to the mobile device's location, for instance, to prevent delays when cloud services are accessed.
- Scientific computing: The use of cloud services on the application layer for scientific computing forms another application category. Since most of the calculations in this environment require enormous computing and data storage capacities for a defined period of time, cloud resources represent a low-cost alternative to grid or cluster systems.

  However, the main problem that arises when scientific applications are implemented is the high latency of communication between the nodes in a cloud computing system, because the majority of such applications were designed to run on dedicated nodes with high-speed networks. The applications need to be adapted on the one hand to enable this high latency to be taken into account in the cloud computing system and on the other, to compensate for variations in system performance.

Figure 4.3 shows the model of the cloud services on the application layer. Applications in domains like finance, procurement, or collaboration are provisioned to cloud consumers as a service. However, the actual look and feel of the runtime environment of an application layer service can vary from one service vendor to another. In Figure 4.3, platform layer services, which in turn are based on infrastructure layer services, are used to run application services. Three different cloud vendors on different layers can have a stake in the service delivery in this kind of scenario.

Runtime environments in which a cloud service vendor on the application layer does not use any other services on another cloud vendor's platform or infrastructure, or where two service vendors – namely an application vendor and a platform vendor – are involved in the service delivery, are possible variants of this model. The runtime environment can impact the security of a cloud service on the application layer, for instance if the cloud vendors concerned apply different security standards and guidelines.

4.4 Platform layer: Platform as a service (PaaS)

The platform layer contains the environment for developing and provisioning cloud applications. The principal users of this layer are developers seeking to develop and run a cloud application for a particular platform. They are supported by the platform operators with an open or proprietary language, a set of essential basic services to facilitate communication, monitoring, or service billing, and various other components, for instance to facilitate startup or ensure an application's scalability and/or elasticity. Distributing the application to the underlying infrastructure is normally the responsibility of the cloud platform operator. The

Figure 4.3: Cloud services on the application layer

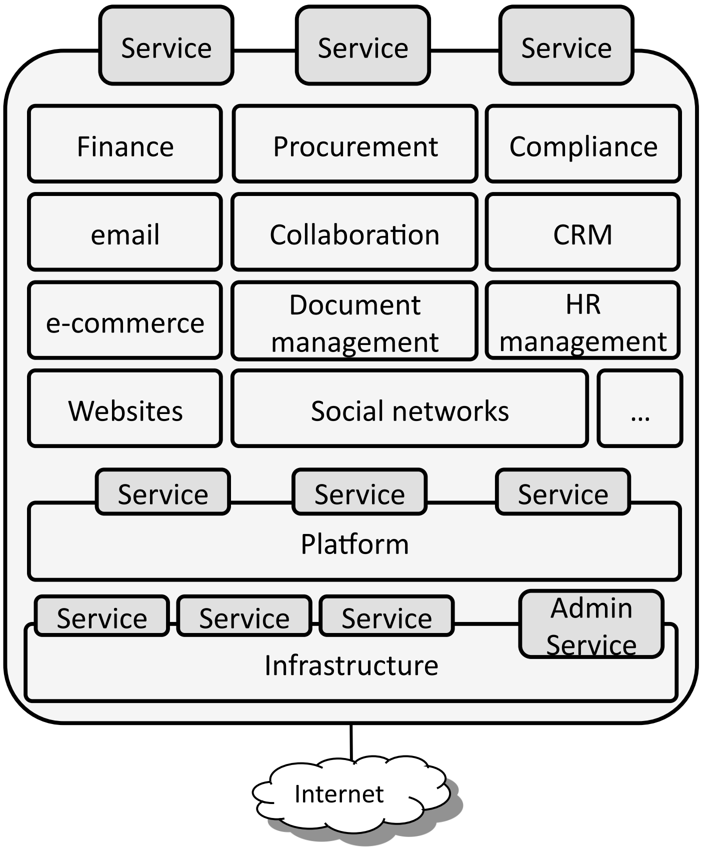

services offered on a cloud platform (refer to Figure 4.4) tend to represent a compromise between complexity and flexibility that allows applications to be implemented quickly and loaded in the cloud without much configuration. Restrictions regarding the programming languages supported, the programming model, the ability to access resources, and persistency are the possible downside.

Since there are no standardized specifications regarding the services and components available on a platform, a single case study must be carried out of the target platform to determine the most suitable platform for a particular application. The term "platform as a service" (PaaS) is often used generically to refer to all services and components offered on the platform layer. Google App Engine[1] and Microsoft Azure Platform[2] are two typical examples of platforms. Both Google App Engine and Microsoft Azure provide basic services that can be used to develop websites, transfer existing applications to a cloud computing system, or design customized applications.

In addition to the above-mentioned cloud service developers, the players on the platform layer are the platform vendors who run the platform and offer its basic services as well as the tool vendors who provide the development tools and/or the programming languages. By integrating legacy platform services

[1] http://code.google.com/intl/de-DE/appengine/

[2] http://www.microsoft.com/azure/default.mspx

Figure 4.4:
Cloud services on the platform layer

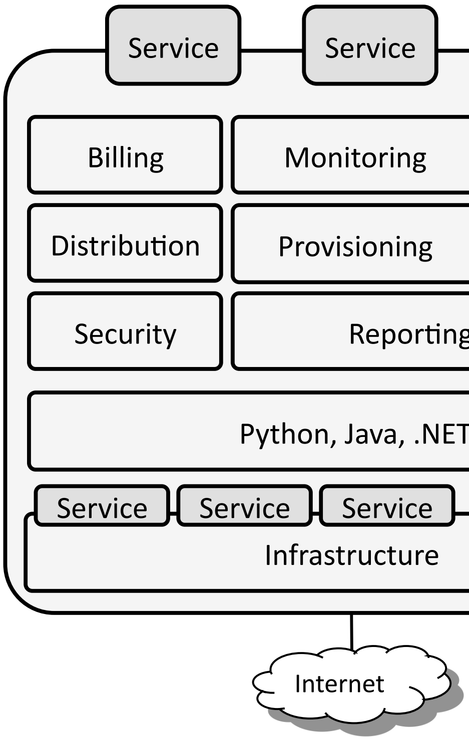

in their applications, developers can reduce the complexity of their software development tasks, accelerate the development process, and consequently limit the number of potential weaknesses by reusing existing source code, because it can be assumed that it is also in the platform vendor's interests to provision basic services of the highest possible quality.

4.5 Infrastructure layer: Infrastructure as a service (IaaS)

The services on the infrastructure layer are used to access essential IT resources that are combined under the heading "infrastructure as a service" (IaaS). These essential IT resources include services linked to computing resources, data storage resources, and the communications channel. They enable existing applications to be provisioned on cloud resources and new services implemented on the higher layers. Three key players can be identified on the infrastructure layer, namely the resource consumer, the resource vendor, and the resource intermediary.

A resource consumer deploys the resources provided by a resource vendor to run applications and provision services on higher layers of the cloud computing system. The resource consumer's job here is to specify the required resources and their service quality in the form of service level agreements that must be complied with by a resource vendor. Resource demand is often automated in a cloud computing system by adding new resources to the existing capacities, so that a constant service quality is guaranteed in spite of load peaks, outages, and other similar events. Since the use of additional resources is also frequently linked to additional costs, resource consumers must define rules to control how certain events impact resource demand. They can do this on the basis of their own experience, reports in forums and journals, best practice approaches (such as those developed by associations like the Cloud Security Alliance), or historical measurements. In a cloud computing system, each vendor's resources are selected and accessed either directly or through an intermediary, who acts as a broker for the resources of different vendors.

Figure 4.5:
Cloud-Services der Infrastrukturschicht

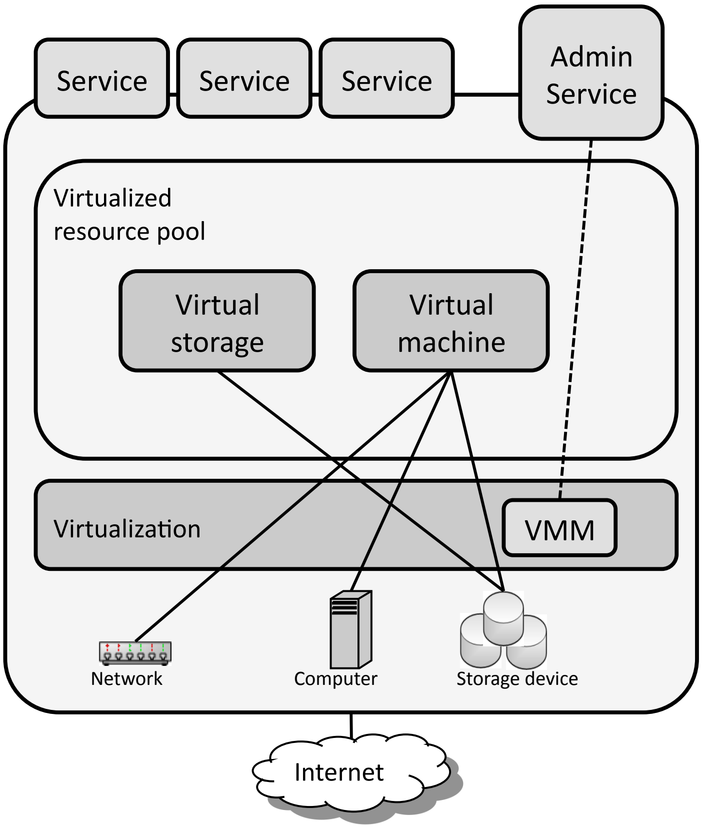

A resource vendor makes virtualized resources available in a cloud computing system. Virtualized resources are computing resources, usually consisting of a CPU and RAM, data storage resources (elastic block store and databases), and network resources (refer to Figure 4.5). These resources are provided by a resource vendor in a standardized way. Individual resources are often combined in resource pools and provided in different amounts. A resource vendor often operates its own datacenter, in which the virtual resources are combined in a resource pool and the physical resources are abstracted using virtualization software. The most popular kind of virtualization software is a VMM (Virtual Machine Monitor), which assigns virtual resources to physical resources and enables virtual resource pools to be created. The physical resources are abstracted by virtualization, which means they can then be shared by several operating systems and end user environments on the virtual resources – ideally, without any mutual interference.

In addition to resource vendors who produce their own resources and offer them to resource consumers, another vendor type can also be identified. This vendors acts as a resource intermediary, playing the role of broker for unused capacities between resource producers and resource consumers. A marketplace or a specialist portal can be used for this purpose.

Computing resources Computing resources are the central resources of a cloud computing system. They are offered as a service as a package of individual resources. Computing resources usually consist of three components, namely a CPU, RAM, and a hard drive, which determine the performance of a virtual machine. The end users of a virtual machine can either select the individual resources in the resource triplet themselves or choose from preconfigured quantities, in other words they define the resource configuration they need to carry out a particular task.

The essential technology that enables computing resources to be offered as a service is implemented using a variety of virtualization technologies, which give vendors the necessary flexibility to configure different virtual machine sizes dynamically while simultaneously protecting their physical infrastructure. Virtualization technologies additionally allow several virtual machines to be isolated on one physical node – something that is useful, for instance, when it comes to implementing the security goal of integrity. In general, the virtualization products are very strongly adapted as a service by the computing resource vendor and not visible to the customer, so that their level of security is extremely difficult to assess. Special protection must be provided for the remote administrative access to the virtual machines, which usually takes place via a public network and frequently includes the option of creating additional virtual machine instances or deleting existing instances – possibly violating one or more of the protection goals.

Data storage resources Storage capacities that allow a player's data and applications to be stored on remote hard drives and distributed data structures (e. g. distributed databases), so that they can be accessed from any location, are likewise essential resources. Data storage resources are offered as cloud services; they make it easy for cloud applications to scale their storage capacity beyond the narrow limits of the virtual machine or server. Data storage services have to satisfy special (security) requirements, for instance if they are used to store end user data or other confidential information. In addition to high availability, reliability, scalability, and data consistency, all the protection goals described in chapter 3 should be complied with.

The techniques and algorithms employed for data storage are specific to cloud services. To simplify processing of large quantities of data, they are normally based on Google's MapReduce framework as well as on other methods for managing data replicas [16]. Free implementations of these methods, such as Hadoop, provide only rudimentary authentication techniques, and could threaten the protection goal of authenticity [12]. The number of replicas, as well as their synchronization and deletion, represents a further potential security problem. When a replica is created, data could cross national borders or transcend the boundaries of organizational structures, leading to privacy violations. End users often have no way of reconstructing the number of replicas, the encryption techniques, or movements of data.

Several security implications can be derived from the methods used to store data, and these must be taken into account when the security guidelines are defined and implemented by the cloud infrastructure operator. The specification of permissible storage locations, data manipulating mechanisms, and long-term data storage are just a few examples here.

Communications channel The importance of the communications channel as a resource is growing as cloud computing systems become increasingly widespread; the communications channel is fast becoming a central component of the cloud infrastructure because no cloud services can be purchased without it. Cloud computing systems must therefore possess a series of capabilities to support dynamic, service oriented infrastructures and guarantee the warranted service quality. Various concepts that attempt to isolate end user communications channels from one another and provide encrypted point-to-point or end-to-end connections are used to achieve these objectives.

In the context of cloud computing, the availability of the communications channel and the ability to ensure confidentiality and integrity comprise one of the most important goals for cloud services. Availability is chiefly threatened by distributed denial-of-service attacks [34], which have been carried out successfully on almost all vendors of cloud computing systems in the past [9]. Most of the possible deterrents to counter distributed denial-of-service attacks in cloud computing systems are designed to detect such attacks as soon as they occur and to limit the resources a cloud service can procure from the cloud vendor's resource pool.

To guarantee the confidentiality and integrity of the communications channel, all messages between two end points should be encrypted prior to transmission. This strategy should also include communications within the cloud as well as with other participating service vendors who are not visible to the end users.

4.6 Cloud service deployment models

Cloud services can be deployed in different ways, depending on the organizational structure and the provisioning location. Three deployment models are usually differentiated, namely public, private (or internal), and hybrid cloud service usage; these are described in detail in the following and shown in Figure 4.6.

Figure 4.6:
Private, public, and hybrid clouds in a cloud deployment model

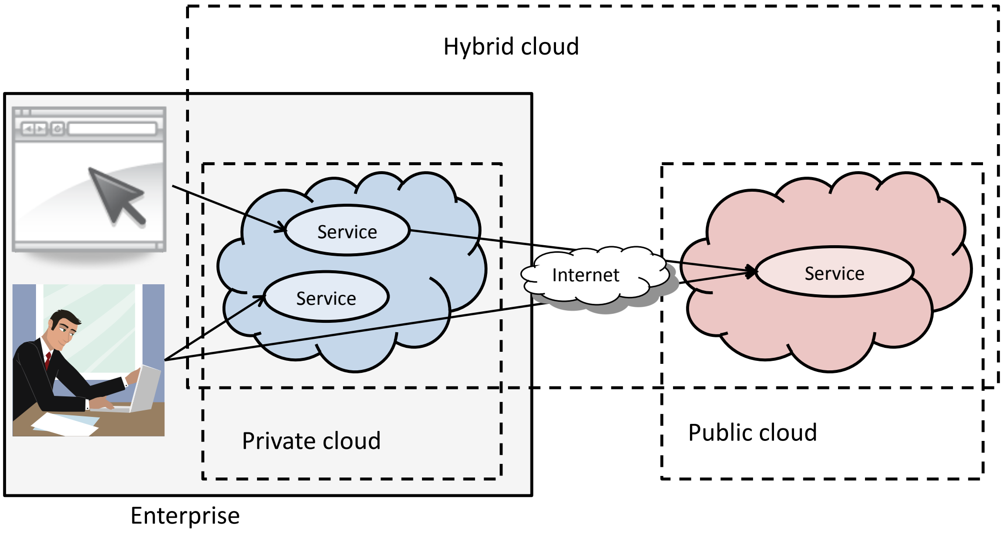

- Public cloud: The deployment of a public cloud computing system is characterized on the one hand by the public availability of the cloud service offering and on the other, by the public network that is used to communicate with the cloud service. The cloud services and cloud resources are procured from very large resource pools that are shared by all end users. These IT factories, which tend to be specifically built for running cloud computing systems, provision the resources in precisely the required quantities. By optimizing operation, support, and maintenance, the cloud vendor can achieve significant economies of scale, leading to low prices for cloud resources. In addition, public cloud portfolios employ techniques for resource optimization; however, these are transparent for end users and represent a potential threat to the security of the system. If a cloud vendor runs several datacenters, for instance, it can assign its resources in such a way that the load is uniformly distributed between all centers. As mentioned earlier, this can result in a loss of consumer control.

- Private or internal cloud: Private (or internal) cloud computing systems emulate public cloud service offerings within enterprise boundaries and within an internal network. Private cloud computing systems make use of virtualization solutions and focus on consolidating distributed IT services. The chief advantage of these systems is that the enterprise retains full control over corporate data, security guidelines, and system performance.

  On the downside, a private cloud also means having to purchase, run, and maintain the IT components – functions that in a public cloud computing system are the vendor’s responsibility. Furthermore, in contrast to public cloud computing systems, important key performance indicators – such as a value of more than 1,000 computers per administrator – can only rarely be achieved, so that the cost savings generated by the private cloud are likely to be much smaller, though at the same time the security risk is much lower.

Private cloud computing systems can serve as reference scenarios for comparing different solutions if a set of scenarios needs to be assessed. However, in view of the high capital costs, most enterprises will not normally set up an internal cloud unless they expect to derive a substantial benefit from consolidating their existing datacenters on a private cloud computing system. For example, economies can be achieved by leveraging the characteristics of cloud computing systems described in chapter 2 as well as through centralized, and largely automatic, system administration.

- Hybrid cloud: A hybrid cloud service deployment model implements the required processes by combining the cloud services of a private cloud computing system and a public cloud. The hybrid model is also suitable for enterprises in which the transition to full outsourcing has already been completed. They can reduce the costs for their existing outsourcing agreements by procuring part of their outsourced services from cheaper cloud vendors.

A common feature of all deployment models is that the cloud services offered are selected by the end users themselves and accessed via the network using web portals or web service technologies. The resources are shared with other end users, possibly with implications for security, and billing is based on resource usage.

Which of the three cloud service deployment models is ultimately chosen by a particular enterprise tends to depend on the security guidelines and legal requirements that have to be complied with. In spite of this, the decision will probably hinge on the process or the application. Critical systems will most likely be transferred to a private cloud computing system by consolidating internal datacenters, while public clouds are more suitable for standardized processes. This strategy results in a hybrid deployment model for cloud services, which in the majority of cases represents a compromise between risks and benefits.

5 Taxonomy of the security aspects of cloud computing systems

Following on from the discussion of protection goals in chapter 3 and the structure of cloud computing systems from a service perspective in chapter 4 this chapter introduces a taxonomy of the security aspects of cloud computing systems. This taxonomy includes the most important security aspects which must be considered whenever cloud services are used. The aim of the taxonomy is to define a flexible framework which, depending on the way a cloud service is deployed, enables decision-makers and IT managers to fine tune security standards in particular areas and which also allows for an overall assessment of cloud security risks. Security risks may be included in a detailed risk-benefit analysis, for example, and drawn on to evaluate suitable security measures.

This chapter discusses the methodology used in analyzing the risks associated with cloud services focusing on the protection goals which ought to be met when using cloud services, structured according to different types of problems which may impair one or several of these goals. However, as the cost of services often plays a very important role in the selection and evaluation of cloud services, financial criteria for the selection of cloud services are also described in addition to the security criteria.

The taxonomy presented here draws on a number of initial studies of cloud security issues prepared by Gartner [24] and the Cloud Security Alliance [10]. The taxonomy in the present study builds on this preliminary work and specifies a more detailed taxonomy of critical security areas in cloud computing. In a sense, a taxonomy of cloud computing risks maps out the security critical aspects involved in procuring cloud services and may be regarded as the starting point for a deeper consideration of security issues.

The four taxonomic categories – infrastructure, application and platform, administration, and compliance and their further subdivision – are introduced in chapter 5.1. The security aspects of infrastructure are described in section 5.2 and the security aspects of the application and platform in section 5.3. The other domains of the taxonomy – administration and compliance – are considered in section 5.4 and section 5.5.

5.1 Taxonomic structure

While the structure of the taxonomy is based on the service model layers of cloud computing systems introduced in chapter 4 it adds two cross-stack factors: administration and compliance. Figure 5.1 illustrates the taxonomic structure. The diagram clearly differs from the familiar three-layer service model. For taxonomic purposes the application and platform layers have been merged, with both layers now being distinguished by lower-level security aspects.

Figure 5.1: Taxonomy of the security aspects of cloud computing systems

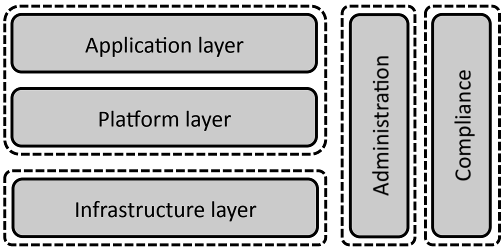

Figure 5.2shows the essential taxonomic structure of the security aspects of cloud computing systems and their lower level security risks. The structure is made up of 4 main domains: infrastructure, application and platform, administration, and compliance. Each of these main areas and the associated security risks are discussed in detail in the following sections with a checklist of specified questions which a cloud consumer would be well advised to ask a cloud vendor.

Figure 5.2: Taxonomy of the security aspects of cloud computing systems

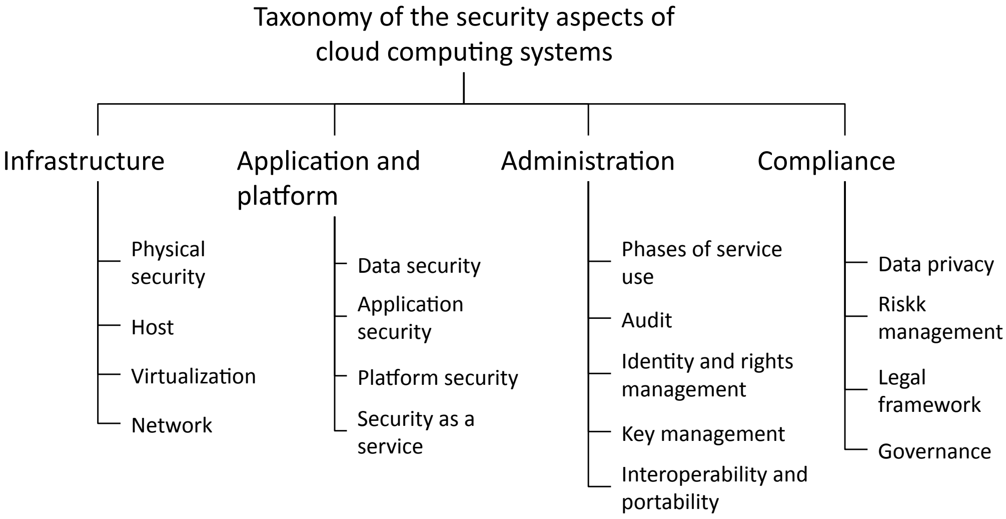

The interdependencies between each of these areas are analyzed in greater detail in the following and shown in Figure 5.3 to provide a further explanation of how the taxonomy can be used. Arrows indicate interdependencies between different areas with the number of inward-pointing arrows showing the importance of each particular area for the overall risk to the security of a cloud service. In this scheme of things, the security of the infrastructure is the most important factor followed by the security of the application / platform domain. The administration area has the lowest level of interdependency, with just one inward-pointing arrow.

The taxonomy integrates two different ways of looking at the security of cloud computing systems. The technical view encompasses the infrastructure, application and platform as well as administration, whereas the process perspective includes compliance. The diagram shows that administration depends on the

Figure 5.3: Interdependencies within each area of the taxonomy

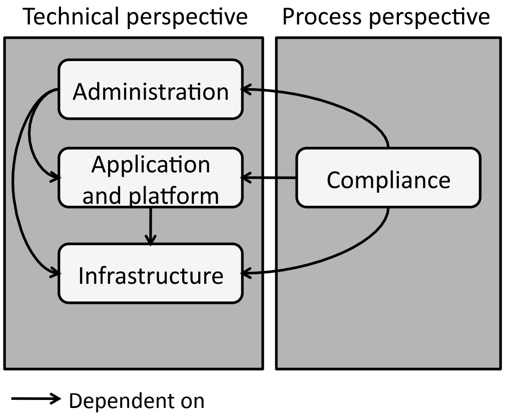

technical properties of the infrastructure, the application and the platform. A cloud consumer can only use the administrative features provided by the cloud vendor. Compliance, on the other hand, is dependent on the technical perspective given that it is extremely difficult to achieve compliance for the cloud consumer without the support of the cloud vendor, its technical systems and organizational processes. This is one reason why it is so important that a cloud consumer not only assesses as many of the security aspects of the taxonomy as possible before selecting a cloud vendor but also agrees with the vendor on measures to reduce the security risks to level defined in advance by the cloud consumer.

5.2 Infrastructure

The infrastructure area of the taxonomy concerns the threats to the security of services on the infrastructure layer. The infrastructure layer is divided into the four areas of physical security, host, virtualization and network which constitute the core components of the cloud infrastructure. Although users of a cloud infrastructure service do not usually have any influence on these core components, they should nonetheless be aware of the potential threats to security which exist at this level. The complexity of cloud infrastructures also makes it very difficult for users to evaluate their security and leaves them with little choice but to trust the cloud resource provider.

5.2.1 Physical security

The physical security of cloud computing systems encompasses the facilities and building services in which cloud computing systems are located or of which they are a part. Examples of security factors include computer power supplies and cooling systems as well as controlled access to the building, video camera surveillance and the location and structure of datacenters [10] [21]. A power failure, for example, can easily impact the protection goal of availability [39]. The removal – i.e. theft – of computers from the building may, for example, contravene the protection goal of confidentiality. The structure of the datacenter can also influence the expandability of the cloud computing system and lead to bottlenecks in the performance offered by a fast-growing cloud vendor.

Most enterprises contract building security out to external firms [7]. However, users should check who is authorized to enter specific areas of the datacenter and to request such information from the cloud vendor in writing where relevant. It is also important to specify the incidents – such as power or CCTV failures, changes in building access controls or the relocation of the computer to a new datacenter – about which the cloud vendor must alert the cloud consumer.

Physical security checklist

- Does the consumer have access to CCTV data or recordings made by the cloud vendor's access control systems in the event that a notifiable incident occurs?
- Do all the cloud vendor's datacenters use the same physical security standards?
- What physical security measures are taken by the datacenter which holds the consumer's data?
- In what way is the datacenter secured?
- How is access to the building secured?
- Are access cards, biometric procedures, video camera surveillance, building surveillance and the permanent accompaniment of guests in the datacenter guaranteed?
- What alarm systems are used?

5.2.2 Host

The host provides the environment in which the processes and their calculations are carried out. This makes very tough demands on security in terms of the protection of the processed data, the availability of the host and the reliability of the calculations carried out on the host.

Potential threats to the data protection goals usually originate in applications running outside the user environment which may affect data within the user environment. If a potential attacker's application is able to influence local data – where local refers to the same physical computer on which both the attacker and user run their applications – it will also be able to change or destroy such data or make the local environment unusable. Isolation helps to keep and run potential attackers' external applications in a protected environment so that, ideally, malicious applications are not able to leave their assigned environment. The virtualization concept is used in cloud computing systems with the aim of isolating several user environments. Direct access to the host resources is no longer allowed but is controlled by a virtual machine monitor. It should be clear and documented at all times which process or which actor has accessed the host. This makes it easier for the user to check the security of the system.

Another threat to the protection goals is the running of applications by a user. When running an application the host resources must often be assigned at the abstraction level of a virtual machine. This can lead to the 'starvation' of an application [41] [14]. An application is said to be starved if a neighboring or higher-priority application utilizes a large amount of a host's resources and thereby makes it impossible to run another application. The significance of this scenario is highly dependent on the intensity with which applications utilize a host's resources and influence its capacity.

In the past, bottlenecks have occurred in commercial cloud service offerings which have resulted in the starvation of applications, in particular due to the overutilization of resources [9]. Bottlenecks have, for example, been caused by distributed denial-of-service attacks intended to impair the reliability and availability of a provider's resources. In order to avoid the risk of starving an application, resource services should be chosen which offer consistently high performance – by requesting monitoring services and analyzing performance history, for example – and/or by imposing contractual penalties for violations of service quality criteria, such as availability.

Host security checklist

- Are procedures adopted which prevent the starving of applications?
- How are the processes of various user applications isolated from each other?
- What procedures are adopted to insulate the host?
- Who has access rights to the hosts in the vendor's datacenter?

5.2.3 Virtualization

As discussed in the previous section, virtualization is mainly used in cloud computing systems to isolate user environments and is consequently an important basic building block in cloud computing systems. At present virtualization is mainly used in datacenters to consolidate computers and to increase the use of the datacenter's capacity. The possibility of using isolation to create a secure environment is a by-product of virtualization solutions and a key requirement for the separation of user environments and compliance with the protection goals defined in chapter 3 .

Threats at the virtualization level often originate in the management of access rights and the dynamic nature of cloud computing systems. Before using cloud computing systems it is important to define very precisely which users should be authorized to administer the virtual machines, how the file permissions for the virtual machine are defined and what authorization the guest operating system has. In addition to the authorization rights relating to the host, authorization must also be defined at the network level, such as the configuration of a host based firewall or access to other Internet or cloud resources.

Current virtualization solutions, such as Xen[1], KVM[2], VMWare ESX[3] or Microsoft's Hyper-V4, offer the migration of virtual machines between hosts, which can violate one or several protection goals, such as a user's privacy. In this connection checks should be carried out to determine whether a vendor of cloud services uses this feature and what the consequences might possibly be. The vendor should also provide information about the geographical location of the virtual machine, or submit a certificate, as this may be stipulated by law.

Virtualization security checklist

- What virtualization technology does the cloud vendor use?
- How does the cloud vendor ensure that the insulation of the virtual machine is complied with and a virtual machine is not, for example, able to access the memory area of others?
- What measures are taken to protect the virtual machines?
- What is done to prevent faulty virtual machines resulting in memory corruption owing to the exploitation of a security hole?
- What is done if a virtual machine monitor (VMM) is compromised?
- How secure is the communications channel between the virtual machines and the VMM?

5.2.4 Network

The network – and its components such as communication protocols and filter technologies – is another important part of the infrastructure which may influence the security of the cloud computing system. The purpose of communication protocols is to enable uniform use to be made of the cloud services by users and between the computers of one or several cloud computing systems, while filter technologies such as firewalls, intrusion detection systems (IDS) or intrusion prevention systems (IPS) enable only certain network connections and in this way prevent malicious intrusions into the system.

[1] http://xen.org/

[2] http://www.linux-kvm.org/page/Main_Page

[3] http://www.vmware.com/de/

The following section summarizes in brief the security aspects of the network of a cloud computing system as a more extensive discussion would go beyond the scope of this study. Cloud computing systems usually only function if they have a reliable network infrastructure; this means that both the cloud user and the cloud vendor need to have an in-depth understanding of network security. The challenges for cloud computing systems from the point of view of network security are based on compliance with the protection goals introduced in chapter 3 and typical requirements for cloud services which should be accessible from anywhere, using any terminal device and using heterogeneous network infrastructures. What is more, the cloud-specific security aspects of networks should also be taken into consideration alongside the secure forwarding of messages and secure multicasting.

Based on the ISO/OSI layer model [1] network access and important security functions can be controlled at various levels, such as at the IP level with IPSec or with TLS/SSL on the transport layer. In this context use is made of procedures for insulating network traffic by means of virtualization, access control by means of firewalls, integration of VPN technologies in cloud services as well as procedures for recognizing and removing suspicious network packages using IDS or IPS systems.

Network security checklist

- What procedures are adopted and network security systems used by a cloud vendor?
- What technologies are used in order to stop network intrusions, such as denial-of-service attacks, man-in-the-middle attacks or port scanning?
- How are these systems configured?
- What configurations can or must the cloud consumer use?
- What response is made to security incidents? Does a process model exist?

5.3 Application and platform

The key risks affecting the application and platform part of the cloud taxonomy are those which can arise during the development and use of cloud services and which may have their origins both in the infrastructure and in the application provided as a service as well as the associated platform. Influences originating with the security of service oriented architectures, and the security of web applications, play an important role in guaranteeing the protection goals for data, applications and processes in cloud computing systems which merit such protection. The following areas have been identified in the framework of the taxonomy of the security aspects of cloud computing systems and are considered in greater detail in the rest of this section: data security, application security, platform security and security as a service.

5.3.1 Data security

In this study data security refers to the security of all data – including any existing configuration and meta data – which is stored and processed in cloud computing systems and transported between cloud computing systems and their services. The focus is on the protection goals of confidentiality and integrity in particular.

In this case a cloud user's data is stored on the cloud vendor's computers in the same way that a company's data is handled when IT services are outsourced to a conventional IT outsourcing provider. This means that the cloud service vendor must implement and provide security functions which protect this data and in some circumstances may be held accountable to the cloud user and asked to specify how data is protected.

Before data is transferred to a cloud vendor's service a cloud consumer should first classify its data and stipulate precisely what data may be stored with a cloud vendor. This classification must specify the exact security measures which must be used to communicate and store data. These may include certain types of cryptographic procedures or guidelines which must be supported by a vendor. One way of ensuring that protection goals are met is to define security guidelines which vendors are required to comply with. Security guidelines may prescribe the use of specific encryption technologies such as public key infrastructures (PKI), for example. The key for the secure transmission and storage of the data is usually exchanged with the cloud vendor and used on this data. Key management is considered in greater depth in the discussion of the management of the taxonomy.

A cloud consumer can also apply the data minimization principle under which customer data is removed from or replaced in the data records processed by a cloud service, for example, only to regain their original semantic meaning when complemented with data kept within the company. This kind of scenario can be used in compute-intensive statistical calculations, for example, in which figures are only required for the calculation in order to assign them to a customer but not for the actual calculation itself.

Data security checklist

- Where is data stored and how is it separated from other customers' data? Is the data on the cloud vendor's computers stored by the vendor in encrypted form?
- Where else is the data stored, e. g. data backups and archiving or using a redundant cloud computing system?

- Who can see the data when it is being stored, during processing and when it is being transmitted through a network?
- Who can access the data when it is being stored, processed or transmitted?
- Is the data secured in a way which restricts its visibility and usability to the data owner?
- If data is deleted, is it also deleted from all application instances, all caches and all backup copies?
- What encryption procedures does the cloud vendor offer? Is the use of these procedures stipulated by contract?
- Can backup copies be encrypted?
- What guidelines and procedures does the cloud vendor use to create, distribute and manage the data and data replicas?
- Can data be recovered after it has been deleted?
- Is it possible to retrieve data in the company again?
- Can stored data be shared with other cloud consumers?

5.3.2 Application security

Application security includes methods and procedures for ensuring authenticated access to cloud services and consideration of security criteria in the development of cloud services. In addition, compliance with integrity, availability and authenticity are also required in addition to confidentiality. As applications in cloud computing systems are usually provided via a public network, the following issues need to be taken into account with regard to application security:

- Messages: Messages should be transmitted in encrypted form using web service security standards (such as WS-Security [2]) and protected against repeated transmission of the same message (replay attack). It should also be possible to sign and validate messages against a schema – e. g. using XML schema standards [3] – in order to be able to identify the sender and to identify incorrect messages before actual processing begins.
- Session: The vendors of cloud computing systems often log sessions themselves for billing purposes. A session is defined as the time between setting up and tearing down the connection to the cloud vendor. In this context it is important to prevent abuse in the form of the malicious hijacking of an inactive session.
- Configuration: A cloud service should be protected against malicious configuration changes. This could, for example, take the form of a special administration interface which is only accessible to a handful of users.

- Exceptions: Exceptions in cloud services should not hinder the use of services by other users.

Other threats to application security may originate in malware infections, media discontinuities when data is processed by the application, man-in-the-middle attacks or from impersonation of cloud users. As cloud services are often web applications the most important threats to web applications are also extremely relevant to cloud computing systems. In this context reference should be made to the work undertaken by the OWASP Foundation, which produces regular summary reports on the most important threats to web applications [36].

These attacks aim to cause damage by violating the protection goals which can, in turn, result in further financial losses from excessive use of a cloud service which is billed to the cloud user. This scenario – in contrast to distributed denial-of-service attacks which are primarily designed to impair availability – is often referred to as a financial attack.

Application security checklist

- What procedures are used in order to separate the application from the data, the platform and the infrastructure?
- What does the cloud service runtime environment look like? What other services are running at the same time? What kind of security functions do they have?
- Are regular security checks carried out on the cloud services by the vendor and external service providers?
- Are the results of security checks documented?
- What authentication mechanisms are offered and are these mechanisms appropriate to the sensitivity of the data?
- What profile and password controls are used in order to avoid abuse?
- Are monitoring tools which can be used to identify security-relevant incidents offered at the application level?
- How are session time-outs dealt with?

5.3.3 Plattformsicherheit

Platform security is mainly of interest to developers of cloud services who use a cloud platform such as Microsoft Azure, Google App Engine or Force.com to develop their own cloud applications. Cloud consumers whose application layer cloud services are run on a cloud platform may benefit from the platform’s security functions or may be affected by a security threat.

Important platform security features relate to the development processes of cloud applications and the tools used in the process. It is also essential that applications and data are segregated on a cloud platform to ensure compliance with protection goals. Secure development processes must be used to ensure such segregation on the platform and in the application. Novel forms of attack based on side channel attacks are a threat in this context. Although side channel attacks are known as methods of attacking hardware security modules, they have also been used successfully in cloud computing systems to circumvent the separation of user environments [33].

Platform security checklist

- What secure development processes are used?
- Are the services provided by a cloud platform subject to ongoing security checks?
- Are the results of security checks documented?
- What measures are taken to isolate applications and data on a cloud platform?
- Where are the user data and applications stored via the cloud platform in the cloud infrastructure?
- How are applications provided on and removed from the cloud platform?
- Does the cloud platform offer functions for complying with protection goals, such as the integrity of information?

5.3.4 Security as a service

Security as a service, which forms part of the architecture of the cloud taxonomy, includes various models which enable value-added security functions paid for by a cloud user to be added to the existing security functions of a cloud service. The objective is to achieve a defined level of security without having to make changes to the service itself. This service can be provided by the cloud service vendor itself or by a trustworthy third party, whereby the cloud resources themselves can be used to provide security services.

Examples in the domain of identity and access management include single sign-on services for cloud vendors or in the domain of the management of cloud services themselves. The management of cloud services, for example, offers functions for the automated management of instances of a service in order to guarantee consistently high availability. If an instance of a service can no longer be accessed, a new instance is provided automatically.

Security as a service checklist

- Is the security architecture documented in full?
- Are special security aspects, such as application and platform security, taken into account, on which the security as a service functions are provided?
- Do the cloud services have a security certificate?
- How can the security functions be integrated as a service? Are there open interfaces and a user-friendly portal?
- Which cloud vendors and services are supported?
- Where is security-relevant data stored?

5.4 Administration

The administration of cloud services presents one of the main challenges from a security perspective. This is still given too little support by cloud vendors, nor are tools available to cloud users – or they are still in the process of being developed – which would enable them to manage their rented cloud services in an integrated and efficient manner. The whole of this domain is still the subject of ongoing research work and it is likely to be some time before the administration of cloud services reaches the same level of quality as has been achieved for other existing programs and tools in corporate networks.

The following sections consider the typical phases of use of cloud services, describe the security aspects involved in the assessment of cloud services, and discuss the security risks involved in identity, rights and key management for cloud computing systems.

5.4.1 Phases of service use

In the taxonomy of the security aspects of cloud computing systems, the phases of cloud service use are regarded in the same way as the transaction steps in e-commerce with the aim of reducing transaction costs. The 5 phases of a transaction are: Initiation, agreement, processing, adaptation and checking as well as implementation. Figure 5.4 shows the service use phases for cloud services.

Figure 5.4: Phases of use of a cloud service

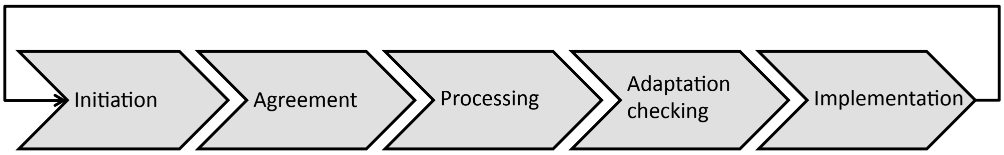

In the initiation phase, the cloud service vendor publishes the description of its product. A cloud consumer launches a search for cloud service offerings with the help of previously defined functional and non-functional criteria. As there is no current standardized search format for cloud services, searches are usually carried out on the websites of the cloud service vendors. In some cases catalogs or directories of a particular platform's cloud services are available which may help cloud users in their search. Cloud vendors themselves rarely provide a description of supported security functions, however, which means that obtaining a rounded picture of a cloud service and its security functions is still a time consuming business. Criteria such as platform independence and interoperability should continue to play an important role to ensure seamless integration in existing IT systems and to enable users to switch from one service vendor to another at reasonable cost [11][10]. At present lack of standardization places strict limits on automated searches for cloud services and this means that the phase of searching for and selecting cloud service vendors inevitably entails high costs.

The initiation phase in the cloud transaction lifecycle is followed by the agreement phase. The next phase, after the cloud user and cloud vendor have come to a satisfactory agreement, is usually the agreement of a contract. Most contracts specify how a cloud service will be used and stipulate rules and duties for both contracting parties. Cloud service vendors offer standardized contracts or service level agreements (SLA) which the cloud consumer only has the choice of accepting or turning down.

It is not usually possible for consumers to negotiate individual contractual terms for service level agreements. Service vendors frequently offer tailored contractual agreements which specify legal aspects in greater detail to their key accounts. The standardized contracts offered by cloud vendors reduce transaction costs and enable consumers to begin using a cloud service more quickly. None of the known service level agreements offered by cloud vendors guarantee specific protection goals; in fact, vendors often refuse guarantees which may arise through the use of third party software products, for example, and on the reliability of which they have very little or no influence at all [29]. During the agreement phase criteria such as the reputation of the product, its user friendliness, previous experiences or knowledge about the technologies used should all be considered when selecting a cloud service to ensure that the service actually chosen is of high quality and is provided by a reliable vendor.

The processing phase focuses on performance of the contract. This is the phase in which the resources are provided, the application is launched, data is transmitted for use, calculations are carried out and results are stored. The specific security risks relevant at this stage have already been discussed in relation to the infrastructure and the application and platform domains of the taxonomy. During the phases of cloud service use other systems for monitoring and measuring service quality and the security functions must be in operation so that, ideally, all the protection goals can be evaluated and the data used to bill for the actual resources consumed. Almost all vendors operate systems which log user actions and the user has little choice but to trust the efficacy of these systems completely for wont of any other means of having the relevant data checked by trustworthy third parties, for example.

After the cloud service has been implemented by the vendor, the cloud resources are dynamically modified to service performance during the adaptation phase in order to achieve the agreed service quality. This involves checking the outcomes of the processing phase as well as making other adaptations for future transactions based on current and past transactions, e. g. the quantity of resources consumed. If the values measured deviate from the agreed performance, dynamic countermeasures for more computing capacity or bandwidth can be defined and introduced automatically. This may entail scaling resources up or down as the case may be.

The implementation phase is the last phase of service use. During this phase deviations from the performance and security metrics agreed in the SLA are analyzed and dealt with, i.e. by payment of penalties for failures to provide the agreed performance or the submission of service vendor ratings which affect the vendor's reputation. These measures may be triggered from the security perspective by the failure of particular security controls which breach a cloud consumer's defined requirements or by too many open security holes in the services offered by the cloud vendor which are not closed within the agreed timeframe.

Phases of service use checklist

- How does the cloud vendor meet security requirements?
- Can the applications and data be recovered if contractual negotiations are broken off by the cloud consumer or vendor?
- Are all the vendor's security functions documented? Is enough information available for assessment purposes?
- What guarantees does the vendor's standardized service level agreement provide? What exceptions are there?
- Does the vendor offer to agree individual service level agreements?
- Where and with which system components is the performance of a cloud service measured? Is it possible to integrate third party services?
- Are dynamic adaptations to resources made by the cloud vendor during runtime or only between the performance of two services?
- What contractual penalties are included in standardized SLAs? When and to what degree are such penalties due?
- What options for implementing the SLAs does the cloud vendor offer? Does the cloud consumer have to notify a breach of contract itself?

5.4.2 Audit

Audits concern the way in which security-relevant events can be recorded, monitored and evaluated in cloud computing systems. This is particularly important in cloud computing systems and it should be possible to audit all the protection goals. The aim of the audit is to secure evidence based on the recorded data. This means that, to ensure that the audit is as comprehensive as possible, evidence must be secured in all the relevant components of a cloud computing system.

As is the case with most of the security aspects of the taxonomy, there is still no such thing as a standardized approach to auditing. However, a generic approach may be adopted based on the inspection of contractually agreed audit trails. Auditing of compliance with protection goals may include existing processes and procedures which are assessed on the basis of contractually agreed documentation duties which may, for example, require a cloud service vendor to carry out regular security checks. These can be undertaken either manually or with IT support.

Given the extremely complex nature of cloud computing systems it can be very costly carrying out such assessments and companies are therefore unlikely to undertake them very often and only in suspected cases of noncompliance. Security checks, e. g. by means of port scanning, can be difficult to carry out, however, as cloud service vendors usually deploy defense mechanisms against what are usually malicious attacks. This is one reason why detailed security check measures should be stipulated when a contract is concluded.

Audit checklist

- What audit options does the vendor offer?
- Is the measurement data made available to the cloud consumer?
- Are regular security checks performed by the vendor itself as well as by external service providers? Are these stipulated by contract?
- To what standards are checks carried out?
- Is it possible for the cloud consumer to carry out its own security checks?

5.4.3 Identity and rights management

An extremely important security aspect of the administrative domain of the taxonomy is identity and rights management which plays a central role in the integration of cloud services in existing IT landscapes. Here the focus is on two attributes of identity and rights management in particular: the ability to adapt existing systems to cloud computing systems to achieve the security goals of authenticity, integrity and confidentiality, on the one hand, and the ability to protect the privacy of cloud users, on the other.

The adaptation of existing access management systems must take account of the characteristics of cloud systems. Whenever cloud services are obtained via a public network the authentication procedure is exposed to threats on the Internet which need to be taken into account when verifying a user and which were briefly touched upon in the section above on application security. What is more, it is not just an enterprise's employees but also its customers and business partners who may need to authenticate themselves to the cloud service. If all the players involved use different identity management technologies, federated identities offer a way of providing authentication between different technology platforms and can therefore take on an important role in the authentication of users in cloud computing systems [31]. Federated identity management involves the distributed storage of the user's identity or of identity attributes on computers on the Internet. The advantage of a federated identity is that the user only needs to authenticate himself once and can also use the cloud services offered by a number of different vendors. Federated identity management can be implemented using standardized technologies such as Security Assertion Markup Language (SAML) or open source standards such as OpenID [40] [32].

In federated identity management the attributes of an identity must be exchanged in the authentication procedure. At the same time it is important to protect the privacy of cloud users as well as possible to ensure that user and other confidential data belonging to one identity does not need to be transmitted to the communication partner. This can be prevented by using pseudonyms which only allow service use to be assigned to cloud users via a temporary identity without transmitting identity attributes.

In addition to identity and access management features, which enable the simple use of cloud services, new approaches are also needed in the domain of administering these systems in order to be able to efficiently handle complex structures with a series of identity providers and a large number of different rights. The focus in this context is on the management of user profiles which enable communication between user and machine – when an end user uses a cloud service – and between two machines – when two cloud services communicate on a largely automated basis. In this context it is possible that the identity data may also be stored on cloud resources and the security of this information also needs to be guaranteed.

Rights management in cloud computing systems often takes the form of an access control list. The advantages of access control lists are that access rights are easy to manage and that rights, in particular, can be withdrawn easily and efficiently by making the corresponding entries in the access control list. It is also very easy to determine which subjects should have which access rights to a specific object such as a file. On the other hand, it can be very time consuming for users to obtain an overview of their current rights. Another problem associated with the management of rights using access lists is that controlling access for long lists is time consuming and inefficient. Alternatives to rights management in cloud computing systems are offered by systems which provide digital rights management of the type often used for multimedia contents and in cloud computing systems to manage access to stored data.

There is currently no standardized process model available in the field of identity and rights management. However, it is imperative that the access path to cloud services is monitored without having to impose limitations on the scalability or dynamism of the cloud system. An initial approach, and one which has already been put into practice, might be to use security proxies representing a central point of access for cloud services which would make it easier to check access rights. The disadvantage of this approach may be that it would restrict the public accessibility of the cloud service.

Identity and rights management checklist

- What identity and rights management standards are supported?
- What rights issue and controlled rights withdrawal processes are used?
- What standards for the provision of identities and user profiles are supported?
- Is the issue of rights transparent?
- Is there a programming interface for the provision and deletion of rights?

5.4.4 Key management

Key management is a core mechanism of the security implemented in a cloud computing system. In this context the complete lifecycle of the key management with the phases key generation, key exchange, key storage, key verification and key destruction must be mapped in cloud computing systems in order to manage trust in a verifiable way between all those involved in a particular cloud computing system.

The fundamental problem is managing a large number of keys for different cryptographic procedures and distributing the keys to participants who were not taken into account when the key management process was originally planned. Building a relationship of trust by exchanging keys between cloud consumers and cloud vendors is made all the more difficult by the ability to meet short-term need for resources in a cloud computing system, which means that these can be procured dynamically from various providers and platforms. Key management must, therefore, be capable of dealing with a number of different key stores and types of key.

Players in a cloud computing system who were not taken into consideration during the planning of key management may be provided a key via an intermediary, for example. However, it is important that the roles of players who carry out encryption and who store keys are strictly separated to ensure that non-authorized access to other keys is prevented. Measures must also be taken to ensure that data is not "taken hostage" by using an unknown key to encrypt data on a cloud computing system and to ensure that data can no longer be accessed in plain text format.

Keys must always be backed up redundantly and protected in the same way as other sensitive data in cloud computing systems; keys must be recoverable in order to avoid the loss of a key. In the worst case this can lead to a scenario in which data can no longer be decrypted and recovered. One potential option in this context is the distributed storage of keys in which part of a key is stored in one key store and another part of the key in another store. Data can then only be accessed if the different parts of the key are reassembled.

Key management checklist

- Does the cloud vendor offer key storage and management services?
- How are keys for the cloud services generated, managed and protected?
- Is the cloud consumer or the cloud vendor responsible for key management?
- Are keys protected against loss?
- How many keys are there for a single user? One or several? Who do the keys belong to?
- Where is data encrypted and decrypted?

5.4.5 Interoperability and portability

The last two security aspects of the administrative domain of the taxonomy concern the interoperability and portability of data and applications in cloud computing systems. The interoperability of cloud computing systems refers to the capability of two or more independent cloud computing systems to work seamlessly together without the need for special agreements between the systems. Interoperability is a criterion which describes support of standards, such as at the interface or protocol level. The platform independence or portability of a cloud computing system, on the other hand, is the property of a cloud service which enables it to run on different cloud computing systems with different service programs at different layers.

Enterprises should bear in mind both interoperability and portability when choosing cloud services in order to avoid lock-in effects and to reduce costs to a minimum when changing to another cloud vendor. Reasons for changing vendor might be to avoid an increase in rental costs, to retrieve data and applications in the enterprise, the discontinuation of a cloud service by the vendor or deterioration in service quality.

Three interoperability and portability scenarios can be distinguished, depending on whether the cloud service is rented by the cloud consumer on the application, platform or the infrastructure layer. When a cloud service is rented on the application layer the cloud service itself belongs to the relevant vendor who processes the data of a cloud consumer. A cloud consumer must be able to migrate data to a new application. For this reason it is important ensure that a cloud consumer always has access to its data and that the data is always in a format which can be processed by the cloud consumer or transformed into a different data format.

On the platform layer the lack of an abstraction layer between application and platform services can mean that significant parts of the source code need to be rewritten during migration. A lack of security functions or recurring platform security risks may, for example, mean that a change of platform is needed. Whatever the case, the data should then be backed up to a second location as it may otherwise be compromised by the security holes in the platform.

Applications which use services on the infrastructure layer are usually run on a virtual machine which enables applications to be copied from one system to another if the same virtualization solution is used. Initial standardization proposals, such as the Open Virtualization Format[4], support simple migration from virtual machines between different systems. Backup copies of virtual machines should be stored in a cloud-independent format and backed up at regular intervals outside the cloud.

A cloud consumer should establish a risk management process in all three scenarios in order to meet the risks which may be encountered during a potential migration. The consumer can use redundancy to reduce dependence on one particular vendor, for example, or can use cloud services provided by different vendors to diversify risk. The relevant details are discussed in greater detail when considering compliance.

Interoperability and portability checklist

- What standards are supported by the cloud vendor to ensure interoperability and portability?
- Is it possible to access data? In what format is data stored?
- Can the stored data be converted into a different format?

[4] http://www.vmware.com/appliances/learn/ovf.html

- Does the cloud vendor's platform support an abstraction layer which supports the porting of an application? Must the cloud consumer take this into account when designing the application?
- Is a vendor's platform compatible with another vendor's platform? What standards are supported?
- What migration options does the cloud vendor offer?
- Will standards and technologies – such as for long-term archival – continue to be supported in the future?
- Are data backup copies stored by the vendor in a vendor-independent format?

5.5 Compliance

The domain compliance brings together all the regulatory issues which may impact the protection goals. The legal framework of data protection laws and legal requirements of companies regarding data storage and processing in cloud computing systems are briefly discussed in the following. A risk management process is also discussed which can be used by cloud consumers to contain the risks involved in using cloud services. Important security guidelines, certificates and standards which a cloud vendor ought to have are also discussed in the context of governance. In general it is the case that compliance monitoring procedures for Internet based services such as cloud services must be extended if they are to cover applications, users and activities in cloud computing systems efficiently.

5.5.1 Data privacy

Information of all kinds – such as word processing documents, videos, customer or financial data – which was previously stored locally or in a corporate network can be stored in a cloud. In fact, a cloud user might even store all its data in a cloud computing system. Data privacy issues may arise whenever a cloud user stores information in a cloud and shares it with other cloud players. The key question which arises in this context is: can information which, by virtue of using a cloud service, is shared with the service provider be stored and processed in a cloud computing system in compliance with current data privacy laws?

Cloud services may be regarded as the commissioned processing of data, as defined by the Federal Data Protection Act (BDSG)[5], where responsibility for

[5] http://bundesrecht.juris.de/bdsg_1990/index.html

processing the data is held by the cloud user[6]. However, this does mean that it must be possible to declare the cloud service provider liable as the contractor if these regulations are breached. This is guaranteed in Europe by the EU Data Protection Directive [15].

The objective is to provide the cloud with technologies which enable enterprises to do business using cloud services and cloud resources while continuing to protect the end user's privacy. This involves the cloud user communicating these privacy requirements to the cloud vendor and checking to ensure that such requirements are complied with. If an enterprise outsources its employee data to a cloud service – as Siemens has done to SaaS provider SuccessFactors[7], for example – the outsourcing enterprise, in its role as cloud user, must ensure that privacy is not violated.

The German Federal Data Protection Act attaches considerable importance to the geographical location at which data is stored. The Act distinguishes between EU countries, countries which provide adequate levels of protection and third countries which do not provide such protection. Data – including personal data – may be stored and processed in EU and other countries which provide adequate protection, while the Federal Data Protection Act prohibits the transfer to and processing of data in third countries which do not offer adequate protection. Under the Federal Data Protection Act this means that cloud services which procure their applications and resources from countries which ensure an adequate level of protection may be used. However, it is important to note that quite different data protection regulations may apply in other countries.

However, if the person affected consents or if it is imperative that data is processed in order to meet contractual obligations, exceptions apply for countries such as the United States of America, Japan or China which are deemed not to ensure an adequate level of protection.

These exceptions apply to cloud services which only operate datacenters in the USA, for example. Enterprises are, for example, entitled to invoke the imperative need to process data to meet their contractual obligations in order, for example, to store and process customer data on Salesforce computers. The problem of remote data processing is exacerbated as there are currently insufficient technical support systems for cloud computing systems for the continuous monitoring of data privacy. The only option remaining to cloud users is to consider the relevant requirements when choosing a trustworthy vendor in the initiation phase, to stipulate privacy requirements in the contractual arrangements and

[6] http://microsite.computerzeitung.de/article.html?art=/articles/2009020/31942144_ha_CZ.html&page=1&ms=/cloud-computing/index.html&pos=4&tpid=ee54f3c7-0de1-40f5-bb23-2cfdf022aee5&pid=ee54f3c7-0de1-40f5-bb23-2cfdf022aee5

[7] http://www.successfactors.de/press-releases/detail/?releaseid=36

to monitor these in the framework of IT governance. In the case of Salesforce, data is handled in accordance with the protection of privacy criteria set down in the Safe Harbor Framework, which is equivalent to a local data protection directive [19].

German and European data privacy regulations both assume in principle that it is possible to determine at any time on what computers and in which datacenters data is physically stored. However, depending on the procedures used, this is not necessarily the case in cloud computing systems. If, for example, a storage service is chosen which stores data in the MapReduce framework, the split data is distributed across several different servers. Although the constituent elements of this data may not contain any personal data at all, this is no longer the case once the data has been reassembled. This makes monitoring difficult, if not impossible.

Data privacy checklist

- What sort of data privacy rules does the cloud service apply? Can the vendor provide a copy of the guidelines in document form?
- To what locations and components of the cloud computing system can data be transferred?
- Does the cloud vendor hold secondary utilization rights to the data?
- Are statistics kept on the data to enable the cloud vendor to optimize the system or carry out market research?
- Who is able to access the data when it is not encrypted, e. g. during processing?
- Who has access to the cloud computing system host?

5.5.2 Legal framework

In addition to data protection legislation which is designed to protect privacy and confidentiality, cloud computing systems can also be restricted by other statutory regulations. Data whose use is restricted by very specific legislation – i.e. health data, the information held by particular groups of professionals such as lawyers or priests, tax information and data held by companies or state organizations which is confined to a particular geographical location or may not be disclosed to third parties – may not, in some cases, be stored or processed on cloud computing systems [13].

The security goal of confidentiality can very quickly be violated and the law broken, particularly if a cloud service vendor is entitled to read, disclose or transfer data. The legal situation may differ starkly from country to country in this respect and it is therefore important to examine the legislation which applies in each case. The entire legal framework is still in its infancy and much exploratory work remains to be done in this context.

Another aspect which falls within this domain of the taxonomy of secure cloud computing is the problem of access to and the processing of data if the cloud service vendor discontinues its service or the service vendor is taken over by another provider which continues running the services originally offered by the previous vendor. This kind of scenario should also be included in the risk management approach discussed above bearing in mind that cloud service offerings will be consolidated in the future. For example, a model might be applied in which data and computers are transferred to a trustee who enables the customer to continue fetching its data while transferring it in an orderly fashion to the new owner's system.

Legal framework checklist

- Does a cloud vendor offer services which are tailored to meet specific statutory requirements?
- Is the cloud vendor authorized to manage particularly sensitive data, such as health-related data? Does the vendor hold the required certificates?
- Can the geographical location be restricted to the extent necessary to meet statutory requirements?
- What liability rules apply when statutory regulations are infringed? For what events is the cloud vendor liable?

5.5.3 Risikomanagement

Cloud service users must implement a process for the management of their cloud vendors which is capable of handling the risks associated with the cloud services being used. When cloud consumers used cloud services they not only outsource a business process or application, but also the risk inherent in the operation of the relevant processes or applications. As events in the past have shown, cloud consumers should be alert to the possibility that cloud services might fail or security risks emerge [9]. Identifying these risks and defining a risk management strategy are therefore important aspects of cloud service use.

Operational risk management is concerned with all the risks which might occur during the ongoing operation of cloud computing systems. Risk management encompasses conventional security as well as procedures for ensuring continuing business operations and disaster recovery mechanisms. Risks are inherent in the operation of cloud services which may, for example, impact the protection goals of confidentiality, integrity or availability. Operational risk management includes all procedures which contribute to dealing with risks from the cloud consumer's point of view.

Cloud consumers can, for example, draw on their own experience, resort to stipulations in guidelines or external services to support their decision making process. The risk should be made explicit in the agreed service level agreement and applicable systems should be available to support risk treatment. A cloud user should be aware that there is always a risk involved in using a cloud service as – given the complexity of cloud computing systems – a service provider can never fully guarantee error-free fulfillment of the service level agreement.

For the purpose of describing the systematic management of cloud risks in the framework of operational risk management a risk management cycle is introduced in the following and assessed in terms of its application to cloud computing systems. The risk management cycle is based on the Risk Management Standard produced by the Federation of European Risk Management Associations (FERMA)[8] and is shown in Figure 5.5:

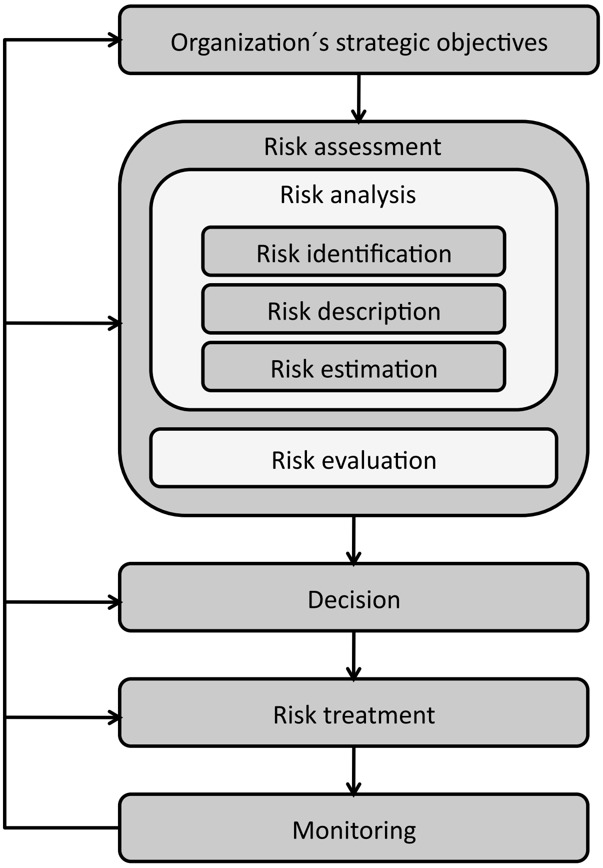

Figure 5.5: Risk management cycle for the use of cloud services

- The organization’s strategic objectives: The reasons for using the cloud service should be derived from the organization’s strategy. Both the cloud consumer and the cloud vendor can decide whether cooperation would serve the organization’s objectives or not by comparing the benefits which the cloud services would provide with the potential risks inherent in such use.

[8] A description of the risk management standard in German can be found at the following Internet address: `http://www.ferma.eu/tabid/195/Default.aspx`

- Risk assessment: Risk assessment includes a risk analysis and evaluation. The risk assessment ends with a decision on whether to accept the risk or to adopt a procedure with which risks can be treated.
- Risk analysis: The first step is to identify the risk in terms of the volatility of the actual service quality and associated security threats.
- Risk evaluation: The risk analysis is followed by an evaluation of the financial impacts of the identified risks based on quantitative indicators which include both the variance as a measure of dispersion of the assessed variables as well as the failure risk as a measure of the loss incurred in the event of failure to fulfill a service level agreement. The failure risk should take account of the fact that cloud users tend to act in a risk adverse manner which means that the failure risk is accorded greater importance than the variance. The risk is evaluated by the cloud consumer who reaches a decision on how to treat risks based on the risk indicators. In the process it is assumed that the indicators have been accurately communicated by the vendor.
- Risk treatment: Risk treatment considers the process of selecting and implementing measures leading to the modification of a risk. The 4 most important risk treatment procedures are risk acceptance, risk transfer, risk avoidance and risk reduction.

  The risk acceptance strategy for the procurement of cloud services would imply all the risks being borne by the cloud consumer and no further efforts being made to modify the existing risks in a contractual relationship between the cloud user and cloud vendor in any way at all. In this case the cloud user usually takes internal risk treatment precautions by holding resources available which will temporarily provide the services in the event of a temporary failure of the cloud vendor's services. The risk acceptance strategy may be used in connection with test systems and proof of concept prototypes for which existing (security) risks usually play a minor role and which may be expected to cause only minor losses for the cloud user.

  The pursuit of a risk avoidance strategy would exclude the use of cloud services for certain types of data and processes as alternative concepts, such as managed services or the provision of services internally, would entail far fewer risks. However, the strategy of risk avoidance is particularly appropriate for important confidential corporate data.

  A risk treatment strategy involving risk reduction would entail a cloud user endeavoring to reduce security risks by, for example, making use of additional security techniques or applying existing procedures at a higher protection level. The objective of risk reduction is to reduce cloud services risks to such an extent that they are acceptable to the cloud user while enabling the latter to still obtain the benefits arising from service use. The application and use of encryption procedures, access management systems, monitoring tools or trustworthy system components is part of general risk reduction and is in this respect probably the most frequently applied strategy when using cloud services.

The last risk treatment strategy to be examined is that of risk transfer. The aim of this strategy is for the cloud user to transfer a certain (security) risk by paying a premium to the cloud service vendor or a third party such as an insurance company. The risk transfer strategy has scarcely been discussed to date in the context of cloud computing systems. However, this strategy does have potential, particularly at the level of infrastructure services. In this context transferring risk may be seem as an alternative way of reducing risk in contrast to the frequently used method of redundancy and can contribute to reducing the costs involved in cloud use.

EAn example for the use of the risk treatment strategy of risk transfer is the "Reserved Instances" product offered by Amazon EC2 cloud services. If a currently operating instance is no longer available it is possible to pay a premium to prereserve another reserved instance which then takes over the service of the failed instance. In this case the cloud user can preempt the risk of failure by paying a premium and prereserving resources which are then available to him when losses are incurred. In contrast to the monetary compensation typically paid in risk transfer transactions in the financial sector, the compensation provided in cloud computing scenarios of this kind is non-monetary in nature. As in the case of compliance with the security goal of availability, other scenarios can also be defined in which it would be appropriate to apply the risk transfer strategy.

- Monitoring: Risk treatment is closely associated with monitoring of contractually agreed service quality on which audits of the services provided by the cloud vendor are based. Service quality or the service level should including monitoring of security metrics, to the extent that they are measurable at all, in addition to time, quantity and utility based values. Grid systems, which are similar in certain respects to cloud computing systems, already employ powerful monitoring systems; the monitoring systems currently used for cloud services, in contrast, measure very few values.

  Another important factor in cloud computing systems is the reliability of monitoring systems bearing in mind that the large number of computing nodes used by cloud vendors' datacenters mean that increased numbers of computer failures and other types of failure are all but inevitable. The aim must be to ensure that monitoring systems are robustly equipped to deal with failures or changes in cloud computing systems. This is also why central monitoring systems are not used in cloud computing systems as they would very quickly run up against their performance limits. Current implementations of monitoring systems take a distributed approach by defining several hierarchical levels which allow monitoring data to be aggregated.

The risk management cycle for cloud computing systems should run parallel to the selection of a cloud service to ensure that threats are identified at the earliest possible stage and to enable risk treatment strategies to be implemented which support achievement of the required protection goals. Modifications must be made to the risk management cycle to fully support the process based on the individual protection goals adopted by a cloud user. In this respect it is important to ensure that the selected risk management strategy is monitored at a later stage by automated IT systems or manually as discussed in greater detail in the next section.

Risk management checklist

- What are the cloud vendor's risk indicators? What are the vendor's objectives?
- What impact could existing risks have on the cloud consumer's business? What residual risks exist?
- How can existing risks be treated? Does the cloud vendor offer services which help to reduce risk, for example?
- What kind of risk management process does the cloud vendor apply? Is this assessed by a third party? Are the documents available for inspection by the cloud consumer?
- Can this process be assessed by the cloud consumer?

5.5.4 Governance

The governance of cloud computing services defines an information security approach in which control systems are established at the process level. This includes the definition of responsibilities for structures and processes, their compliance with previously defined metrics, the stipulation of information security objectives and associated guidelines and criteria for measuring the effectiveness of information security processes.

The challenge inherent in governance for cloud consumers is to find a compromise between the work involved in creating the processes, collecting the data and implementing the processes, on the one hand, and the costs incurred as a result on the other. From a functional point of view the challenge of governance is to define a comprehensive information security framework for cloud computing service procurement and deployment models. Account must be taken of collaboration with a cloud vendor at the information security level, on the one hand, and the responsibilities for implementing and managing security processes and the associated controls between cloud user and cloud vendor defined on the other.

Information security guidelines may be based on a generally recognized standard such as the IT Baseline Protection Manual issued by the German Federal Office for Security in Information Technology (BSI), publications issued by the European Network and Information Security Agency (ENISA) or various security guidelines issued by the U.S. National Institute of Standards and Technology (NIST). A number of different certifications and standards are also offered which define rules for handling data, administrator rights, statutory provisions and other IT security processes. These certificates are usually issued by external organizations and audited on a regular basis.

Examples include the Statement on Auditing Standards (SAS) Number 70 Type II Certificate and the ISO/IEC 27001:2005 Certificate. A service organization's control activities and objects, including the control of information technologies, are documented and confirmed by an external auditor in the SAS 70 Certificate. As one particular type of service organization, cloud service vendors can use the SAS 70 certificate to signalize to potential cloud consumers that they have installed appropriate control systems for their IT-related technologies and processes.

ISO/IEC 27001 specifies requirements for the establishment, implementation, control, updating and improvement of a system for the management of an enterprise's security risks. It does not mandate specific security mechanisms, but is restricted solely to the management level. An ISO/IEC 27001 system includes various plan-do-check-act cycles which result in security mechanisms being subject to an ongoing process of assessment and modification which enable them to keep apace of changes in threats to and weak spots in IT systems as well as the influence of threats to IT operations. With regard to cloud computing systems, management cycles should be modified to the security threats arising from the use of cloud computing systems and should also take account of the security fields described in the taxonomy.

Governance checklist

- Who is liable in the event that data is lost or misused?
- Who holds what part of the data and applications?
- Who is responsible for network access management, reporting, change management, development and maintenance?
- What controls exist on the application, platform and infrastructure layers?
- What certificates does the cloud vendor hold? What security aspects are covered by these certificates?
- Are certificates only audited internally, or are they audited externally as well?
- Are copies of the certifications issued?

- How often is certification performed?
- Who is responsible for security failures? Who is responsible for doing what?
- What processes does the cloud vendor use with its suppliers? Is security a criteria in the selection of suppliers?
- What security controls are applied when using third party software?
- What happens if a cloud vendor's services are no longer available after a business enterprise has been wound up, for example? Is the cloud vendor dependent on the effects of external services which may have an influence on cloud consumers?
- Are the vendor's processes consistent and complete?

5.6 Summary

The taxonomy of the security aspects of cloud computing systems defines – from the point of view of a cloud consumer – a comprehensive framework for assessing the security risks involved in the use of cloud services. The taxonomy consists of four main domains: infrastructure, application and platform, administration and compliance. Important security aspects which may impact the security of cloud services have been defined for each of these areas.

The infrastructure domain focuses on secure datacenter operation and the security of the cloud services offered on the infrastructure layer (e. g. processing). The security of the application and platform as well as services which offer security functions for cloud computing systems are the key issues for the application and platform domain. These services are usually provided by third parties. Management tasks which are essential for the secure use of cloud services are considered in the context of administration. The last domain of the taxonomy is compliance, which presents important security aspects from a process perspective.

Given the mutually interdependent nature of these various domains it is important to consider security as a whole. Checklists for each security aspect of the taxonomy have been drawn up for this purpose and should be presented to a cloud vendor by a cloud consumer. These checklists are used in the next chapter in connection with the example of Amazon EC2 cloud services and the results presented.

6 Cloud services and their security functions

There are a number of cloud services which can be classified according to various service groups, such as infrastructure, platform, application, administrative services and security as a service. The composition of cloud computing vendors and services changes on a daily basis as new providers and services come and go. Initially a selection of cloud services from various service groups and their costs[1] is presented followed by an examination in section 6.2 of providers in terms of their security functions. Reference is made in this context to architecture, infrastructure, administration and compliance. Section 6.3 applies the taxonomy to the cloud vendor Amazon and conclusions are outlined in section 6.4. The websites of the cloud vendors and their white papers are used as information sources.

6.1 Market overview of important vendors

Cloud vendors offer a diverse array of very different services. Offerings differ both in terms of their functionality and the hardware and payment models which they use. The hardware offered by various providers differs, for example, in terms of CPU and RAM size. The payment models operated by cloud vendors differ likewise. Vendors may, for example, offer their services under a pay-as-you-go model, may demand a one-time fee in combination with additional usage charges calculated according to time of use, or ask a fixed price for a service.

A selection of cloud services in the fields of infrastructure, platform, application and administrative services as well as security as a service are presented in the following. This selection takes account of cloud vendors which have been established on the market for a considerable time and/or have a large customer base.

6.1.1 Infrastructure services

Infrastructure services include the provision of computing capacity, data storage and databases. This section considers the computing capacity and data stor-

[1] July 2009

age offered by the five infrastructure providers Amazon[2], Microsoft[3], GoGrid[4], FlexiScale[5] and Rackspace[6]. The provision of databases is considered in relation to the prices offered by Amazon and Microsoft. Prices for the transfer of data arriving at and leaving datacenters are looked at separately, as some providers demand different prices for inbound and outbound data.

In the following, infrastructure services are grouped according to the renting of computing capacity, data storage and the database.

Computing capacity

Amazon is the only vendor referred to here to offer different prices for Europe and the USA as well as two quite different price models. The infrastructure can be paid for on a pay-as-you-go basis or as reserved instances. Under the reserved instance option a one-time payment is made to use the infrastructure for a period of one or three years at a much cheaper rate than under a pure pay-as-you-go model. Each time a reserved instance is run, an additional usage rate is charged for the actual run time.

As it is very difficult to compare the different price models directly, the following focuses solely on the computing capacity of vendors which operate a pay-as-you-go model. Different cloud vendors, and indeed in some cases single cloud vendors themselves, offer variously scaled CPU resources and memory (RAM). As a complete breakdown of all the instances offered by each of the vendors would exceed the bounds of this study, the prices demanded by most vendors are presented with a price tier for the different types of instances. The prices in the lower tier are for smaller instances with small CPUs and little memory; the higher prices are for instances with larger CPUs and more RAM. Microsoft is the only vendor to quote a fixed price for an instance.

The prices for the vendors' server instances are shown in table 6.1 and the prices for data transferred in and out in table 6.2. In some cases the quoted prices for data transfer vary as the price is dependent on the quantity of data transferred.

Table 6.1: Prices for server instances

Vendor	Server instances
Amazon EC2	$0,11 - $1,28 /hour
Microsoft Windows Azure	$0,12 /hour
GoGrid	$0,095 - $1,32 /hour
FlexiScale	$0,04 - $0,64 /hour
Rackspace Cloud Server	$0,015 - $0,96 /hour

[2] http://aws.amazon.com/

[3] http://www.microsoft.com/windowsazure/

[4] http://www.gogrid.com/

[5] http://www.flexiscale.com/

[6] http://www.rackspacecloud.com/?RCMP=cleanEntry

Table 6.2: Prices for data transfer

Vendor	Data transfer inbound	Data transfer outbound
Amazon EC2	$0,10/GB	$0,10 - $0,17/GB
Microsoft Windows Azure	$0,10/GB	$0,15/GB
GoGrid	$0/GB	$0,50/GB
FlexiScale	$0,12/GB	$0,13 - $0,17/GB
Rackspace Cloud Server	$0,08/GB	$0,22/GB

From the lowest prices in the price tier it is apparent that the cheapest cloud servers are those provided by FlexiScale and Rackspace. The highest prices are charged by Amazon EC2 and GoGrid, although it is important to bear in mind that these prices refer to various instances with different CPUs and RAMs. As there is no price range for inbound data transfers these prices are very easy to compare. A range of quantity-determined prices do, however, apply to outbound data transferred by Amazon and FlexiScale. Nonetheless, the prices for inbound and outbound data are very similar with the exception of GoGrid which makes no charge at all for inbound data but charges much higher prices than any of the other vendors for outbound transfers.

The choice of cloud vendor will depend on the cloud user's specific situation and required resources; the same vendor is not equally appropriate for every consumer.

The following example calculations demonstrate the prices charged by the cloud vendors considered so far for hosting small, medium and large-scale websites. The prices are made up of charges for computing capacity and inbound and outbound data transfers. Monthly prices are calculated in each case based on the use of an assumed 732 hours of computing capacity per month. Table 6.3 shows the volume of inbound and outbound data transfers. The prices for hosting a small website are stated in table 6.4, for a medium-sized website in table 6.5 and for a large website in table 6.6.

Table 6.3: Volumes of inbound and outbound data

Website	Data transfer inbound	Data transfer outbound
Small	1 GB	2 GB
Medium	12 GB	120 GB
Large	90 GB	900 GB

The figures in the tables show that no single provider offers the most cost effective resources to meet every different kind of requirement. In this example, Rackspace offers the lowest priced resources for a small and medium-sized website, while FlexiScale offers the cheapest solution for a large website. This demonstrates that cloud vendors must always be evaluated in relation to the requirements of the cloud user. It may, for example, be more cost effective to

Table 6.4: Example for the hosting of a small website

Vendor	Computing capacity	Data transfer inbound	Data transfer outbound	Price
Amazon	$80,52	$0,10	$0,34	$80,96
Microsoft	$87,84	$0,10	$0,30	$88,24
GoGrid	$69,54	$0	$1,00	$70,54
FlexiScale	$29,28	$0,12	$0,34	$29,74
Rackspace	$10,98	$0,08	$0,44	$11,50

Table 6.5: Example for the hosting of a medium-sized website

Vendor	Computing capacity	Data transfer inbound	Data transfer outbound	Price
Amazon	$80,52	$1,20	$20,40	$102,12
Microsoft	$87,84	$1,20	$18,00	$107,04
GoGrid	$69,54	$0	$60,00	$129,54
FlexiScale	$29,28	$1,44	$20,40	$51,12
Rackspace	$10,98	$0,96	$26,40	$38,34

obtain additional data storage for a website from a quite different cloud vendor.

Database

The prices for data storage charged by Amazon and FlexiScale are scaled according to the volume stored and both vendors therefore quote a range of tiered storage prices. Microsoft, Rackspace and GoGrid charge the same prices for every GB of storage capacity, although the first 10 GB/month are provided free of charge by GoGrid. As is the case for data transfers in the computing capacity section, a price tier is quoted for outbound transfers, depending again on the volume of data handled. Table 6.7 shows the data storage prices charged by each vendor and the prices for requests, such as PUT, POST and LIST. The request prices charged by Rackspace Cloud Files are determined by the size of the request file: files under 250 KB are free, whereas files over 250 KB cost $0.01 /month for 500 requests. Table 6.2 shows the prices for data transfer which are the same as the prices for computing capacity data transfers.

Table 6.7 shows that storage prices are very similar, with the exception of FlexiScale which charges significantly higher prices than the other cloud vendors. Amazon and FlexiScale prices are dependent on the volume of stored data while Microsoft, GoGrid and Rackspace offer fixed prices for unlimited file storage. GoGrid is the only one of these vendors to offer 10 GB of cloud storage per month free with every account.

Among all the infrastructure vendors considered here GoGrid offers the lowest storage prices given that it too offers free requests. However, as is the case

Table 6.6:
Example for the hosting of a large website

Vendor	Computing capacity	Data transfer inbound	Data transfer outbound	Price
Amazon	$80,52	$9,00	$153,00	$242,52
Microsoft	$87,84	$9,00	$135,00	$231,84
GoGrid	$69,54	$0	$450,00	$519,54
FlexiScale	$29,28	$10,800	$153,00	$193,08
Rackspace	$10,98	$7,200	$198,00	$216,18

Table 6.7:
Data storage prices

Vendor	Storage	Requests
Amazon S3	$0,15 - $0,18/GB/month	$0,12 per 1000 requests
Microsoft Windows Azure	$0,15/GB/month	$0,01 per 10000 requests
GoGrid	10 GB/month free, thereafter $0,15/GB/month	$0/GB
FlexiScale	$0,43-$0,49/GB/month	$0/GB
Rackspace Cloud Files	$0,15/GB/month	$0 - $0,01 per 500 requests per month[0,5em]

with computing capacity, the best solution depends on the cloud user's specific requirements and the services which the vendor makes available.

Database

Two types of database service are considered: Amazon's SimpleDB and Microsoft's SQL Azure service. Both services operate quite different payment models. While Amazon charges a fee per GB and month for SimpleDB services and an extra fee for machine hours consumed, users of Microsoft's SQL Azure can choose between two editions for which a fixed price is paid for a defined amount of database. The Web Edition of SQL Azure offers up to 1 GB and the Business Edition up to 10 GB of database. Database prices and the required machine hours per request are detailed in table 6.8. Data transfer prices are shown in table 6.9. The price charged by Amazon for outbound data is determined in its turn by the volume of data transferred and prices are tiered accordingly.

Table 6.8:
Database prices

Vendor	Storage	Machine hours
Amazon SimpleDB	1 GB/month free, thereafter $0.25/GB	25 machine hours ree, thereafter $0.14/hour
Microsoft SQL Azure	Web Edition up to 1 GB $9.99/month Business Edition up to 10 GB $99,99/month	$0/hour $0/hour

Table 6.9: Data transfer prices

Vendor	Data transfer inbound	Data transfer outbound
Amazon SimpleDB	1 GB/monath free, thereafter $0,10/GB	$0,10 - $0,17/GB
Microsoft SQL Azure	$0,10/GB	$0,15/GB

The two database services operate different payment models. Amazon's SimpleDB service offers a pay-as-you-go model while Microsoft, in contrast, offers metered units of database storage. Amazon also charges for the machine hours used to complete a request, while this is included in the pack price offered by Microsoft. The database user must choose the service most suited to its preferred payment model and database requirements.

6.1.2 Platform services

Platform services make platform oriented resources and IT infrastructure available for the development and provision of cloud applications. The following section considers the services offered by Google[7], LongJump[8] and Force.com[9]. These three vendors all provide an application development and hosting platform for the creation and hosting of user's own web applications. Google offers the application programming languages Python and Java. LongJump enables the creation of applications with Java and JavaScript. A plug-in can also be used to develop applications directly in Eclipse. Force.com enables both previously developed applications to be used with its point and click functionality and users' own applications to be developed with the Java-like programming language Apex.

Force.com and LongJump bill for complete packages of services which differ in terms of factors such as the number of prebuilt objects and available storage. The user fee, number of prebuilt objects, and the size of the available storage offered by both Force.com and LongJump are shown in table 6.10. LongJump distinguishes between data and document storage. Two types of storage are shown together in the table with the data storage always being equal to $\frac{1}{5}$ of the document storage. In other words, the Bronze Edition provides 5 MB of data storage and 25 MB of document storage, the Silver Edition 10 MB of data storage and 50 MB of document storage and the Gold Edition 20 MB of data and 100 MB of document storage. LongJump offers additional storage if the available space is insufficient. Additional data storage of 50 MB costs $49/month and 250 MB of document storage costs $49/month.

[7] http://code.google.com/intl/de-DE/appengine/
[8] http://longjump.com/index.htm
[9] http://www.salesforce.com/platform/cloud-platform/

Table 6.10:
Prices for Force.com and LongJump platform services

Vendor	Fee	Object	Storage
Force.com			
Free	Free	10	10 MB per user
Enterprise	$50 per user/month	200	20 MB per user
Unlimited	$75 per user/month	2000	120 MB per user
LongJump			
Bronze	$30 per user/month	10	30 MB
Silber	$60 per user/month	200	60 MB
Gold	$90 per user/month	2000	120 MB

Google App Engine, on the other hand, offers its service free up to certain quotas. The quota and the prices for services used in excess of the fee threshold for storage and Google App Engine CPU time are shown in table 6.11 and for data transfer in table 6.12.

Table 6.11:
Google App Engine prices

Vendor	Storage	CPU time
Google App Engine	1 GB/Tag frei, dann $0,15/GB	6,5 CPU-Stunden/Tag frei, dann $0,10/CPUStunde

Table 6.12:
Prices for Google App Engine data transfer

Vendor	Data transfer inbound	Data transfer outbound
Google App Engine	1 GB/day free, thereafter $0,10/GB	1 GB/day free, thereafter $0,12/GB

It is not possible to generalize about which of the platform services considered here is the right one for a cloud user. While they all offer application development and hosting platforms, the vendors considered here differ immediately as far as the provision of programming languages is concerned, although Java, .Net or Python are used and supported in most cases. While Google provides its services free up to a certain quota and then on a pay-as-you-go basis, Force.com and LongJump offer packages prices which differ according to factors such as the amount of available storage and the number of prebuilt objects.

The "right" vendor depends on the cloud user’s individual requirements.

6.1.3 Application services

There is now a huge choice of application services on offer in the Internet with an array of email applications and spreadsheet or CRM (customer relationship management) applications available for rent. This section provides a brief overview of the application services offered by Google Apps[10], IBM Lotus Live[11], Microsoft Office Live[12] and Salesforce CRM[13].

Google Apps

The Google Apps application suite provides the following applications:

- Gmail
- Google Calendar
- Google Talk (instant messaging and voice over Internet protocol)
- Google Docs (word processor, spreadsheet and presentation applications)
- Google Sites (webpage creation)
- Google Video (video hosting and streaming)

Three editions of Google Apps are available: a free advertising-financed standard edition for private users, an education edition for schools and universities, and a premier edition for businesses. The standard and education editions are free. The premier edition costs $50 per user per year.

IBM LotusLive

LotusLive offers applications such as email, Web meeting and social networking applications. The prices for these services vary depending on whether they are provided monthly or yearly, although yearly prices are lower. The prices in the following brief overview of LotusLive Notes, LotusLive Meetings, LotusLive Events and LotusLive Connections services are based on monthly payments

LotusLive Notes is an email, calendaring and scheduling solution. LotusLive Notes costs $9.00/month.

LotusLive Meetings is a Web meetings services which provides a full range of functions for sharing information, giving presentations and demonstrating software. The service costs $48.00/month.

LotusLive Events includes the same services as LotusLive Meetings plus an event management service which can be used to send automatic reminder emails or to access guest registration information, for example. LotusLive Events costs $99/month.

LotusLive Connections is a social networking and collaboration service which costs $12.20/month.

[10] http://www.google.com/apps/intl/en/business/index.html

[11] https://www.lotuslive.com/

[12] http://www.officelive.com/de-DE/

[13] http://www.salesforce.com/de/crm/service.jsp

Microsoft Office Live

Microsoft offers Office Live in two editions: Office Live Workspace and Office Live Small Business. Office Live Workspace provides free online storage and document sharing as well as the Office programs Word, Excel and PowerPoint. However, users can only save up to 5 GB of information.

Office Live Small Business enables users to design their own website and provides web design tools, email accounts and the Office programs Word, Excel and PowerPoint at no charge. The following services in Office Live Small Business Edition are, however, subject to billing:

- Domain name registration costs 9.99/year for .de and .eu addresses or 11.99/year for .com, .org and .net addresses
- Premium email (ad-free email accounts) for 19.03/year
- Additional storage costing between 4.75/year and 14.27/year, depending on storage space
- Additional users at a cost of between 14.27/month and 124.94/month, depending on the number of users

Salesforce

Salesforce.com provides four editions for customer relationship management solutions, each offering a diverse array of functions. The Group edition has the fewest functions. Salesforce offers sales and marketing applications such as account management, contact management, creation of email templates, sending of mass emails or data validation. In addition to these modules, the Professional edition also offers solutions for call center personnel, such as case queues and automatic assignment as well as customizable dashboards. The Enterprise edition offers applications such as territory management, sales and planning management and realtime database mirroring. The Unlimited edition has the largest array of functions and, in addition to the Enterprise edition, also includes automated synchronization of data with a preferred mobile device as well as mobile access to Salesforce applications. Other differences between each of the editions – such as price, number of prebuilt applications, maximum number of supported subscribers per edition and available storage – are shown in tables 6.13 and 6.14.

This is merely a selection of the application services available. The "right" vendor will depend in all cases on a user’s individual requirements.

Table 6.13: Prices and number of prebuilt applications with Salesforce

Edition	Price	Prebuilt applications
Group Edition	75€/Benutzer/Jahr	1
Professional Edition	840€/Benutzer/Jahr	5
Enterprise Edition	1620€/Benutzer/Jahr	10
Unlimited Edition	3240€/Benutzer/Jahr	Unbegrenzt

Table 6.14: Maximum number of supported subscribers and storage space with Salesforce

Edition	Maximum number of supported subscribers per edition	Storage
Group Edition	5	1 GB total
Professional Edition	Unlimited	20 MB per user
Enterprise Edition	Unlimited	20 MB per user
Unlimited Edition	Unlimited	120 MB per user

6.1.4 Management services

Third party management services can be used to manage the infrastructure or applications. This section takes a brief look at the management services Scalr[14] and RightScale[15]. Scalr can only manage the Amazon EC2 infrastructure. RightScale can manage cloud infrastructures provided by various vendors, such as Amazon and GoGrid.

Scalr

Scalr is a redundant and scalable management service for Amazon EC2. The service enables server farms consisting of EC2 instances to be predefined. If demand for resources increases or if one or several instances fail, new instances are automatically provisioned and decommissioned again when they are no longer needed. A number of base images such as load balancer, application server or databases are available to build server farms with. The management service provided by Scalr costs $50/year and the costs for the EC2 instances must be paid separately to Amazon.

RightScale

RightScale is another management service which deploys and manages applications and infrastructures in the cloud. The advantage of this vendor is that several cloud infrastructures from multiple providers, such as Amazon, FlexiScale, GoGrid or Rackspace, can be managed. The costs for these providers' infrastructures must be paid separately and are not included in the fees charged

[14] https://scalr.net/login.php

[15] http://www.rightscale.com/

by RightScale. RightScale offers several editions ranging from scalable websites through to scalable batch processing. The Website edition offers all that is needed to deploy a scalable website in the cloud. The Grid edition allows users to control and manage grid computing and batch processing worker tasks in a scalable, fault-tolerant environment. The Enterprise edition offers all the Website and Grid edition features. The Premium edition offers multi-cloud support and a higher level of customer service in addition to the administrative features provided by the Enterprise edition. RightScale also offers a Developer edition with which users can test some features at no charge. Table 6.15 compares the different editions and shows the prices for each – the one-time usage fee and a minimum turnover per month.

Table 6.15: Prices for RightScale editions

Edition	One-time usage fee	Minimum turnover
Developer	Free	Free
Website	$2500	$500/month
Grid	$2500	$500/month
Enterprise	$4000	$1000/month
Premium	$10000	$4000/month

This section has only looked at the management services provided by Scalr and RightScale. The major advantage of RightScale – in comparison with Scalr, which only offers the management of Amazon EC2 infrastructures – is the management of several infrastructures. RightScale charges considerably higher prices than Scalr, but also offers different features in each of its editions. Which of these services are the right ones is up to the individual user to decide.

6.1.5 Security as a service

Security services for various applications and vendors are offered by a number of different third party providers. This section takes a brief look at the following three services: Google Message Security[16] email protection, PingIdentity's[17] user management and single sign-on service and CohesiveFT's VPN-Cubed solution for EC2[18] which provides an overlay network for Amazon EC2.

Google Message Security

Google Message Security powered by Postini is a software service which secures inbound and outbound email. Spam, viruses and other email threats are blocked and prevented from reaching the enterprise. Users can configure spam protection settings themselves. Google Message Security enables email encryption with TLS (Transport Layer Security) as well as the enforced encryption of all

[16] http://www.google.com/postini/email.html#archive
[17] http://www.pingidentity.com/
[18] http://www.cohesiveft.com/vpncubed/

communications between designated email domains. Google Message Security Service costs $12/user/year. Postini also offers Google Message Discovery, an archiving service which contains the same features as Google Message Security plus email archiving. This service costs $25/user/year for one year of email archiving and $45/user/year for 10 years of archiving.

PingIdentity

PingIdentity's PingConnect is an on-demand single sign-on (SSO) and account management service. PingConnect supports more than 60 software services, such as Google Apps, Salesforce CRM, Postini (Google) or SuccessFactors. The service costs €1/user per application and month.

VPN-Cubed für EC2

CohesiveFT's VPN-Cubed for EC2 product provides an overlay network for Amazon EC2 which can establish a secure connection in the Amazon environment.

Two variants of VPN-Cubed for EC2 are available. The free variant includes two VPNCubed managers. The VPN-Cubed managers can connect two servers either within a single region (EU or US region) or between the two regions. The second variant costs $0.05/hour and includes 4 VPN-Cubed managers which can be used with four servers, for example, within and/or outside a region.

Security services are also offered by third parties which secure existing rented services, although the available choice is not as great as for applications or infrastructures, for example. Users must decide for themselves which additional security services they need to meet their requirements.

6.2 Security functions offered by current cloud vendors

The taxonomy in chapter 5 will now be applied to selected vendors or services and current security functions considered. It is not possible, however, to provide a full list of all the available services and their current security functions at this point.

Section 6.2.2 begins by examining the way data is protected and encrypted in the cloud. Section 6.2.1 then briefly outlines the physical security aspects of datacenter operations and cloud vendors' network security. Section 6.2.3 looks at the service level agreements offered by cloud vendors and section 6.2.4 outlines the certificates held by various cloud vendors.

6.2.1 Infrastructure

This section deals with the physical security of operational datacenters, with the network security issues already identified in chapter 5.2. as well as with measures to secure datacenter operations and the networks of the following cloud vendors: Amazon [5], Google [4], GoGrid[19] and Microsoft [6].

There are number of issues which need to be considered in connection with datacenter security. These include the site at which a center is set up through to security systems and the access control measures needed to protect the datacenter. The site of the datacenter should not be in an area at risk of flooding or in an earthquake zone. Google's datacenters, for example, are located in areas which are protected as far as is possible against possible disasters.

The datacenter site itself as well as computer and critical infrastructure rooms should be kept under surveillance. Amazon has security guards controlling the sites of its datacenter and monitors access to the building with the aid of video cameras, intruder detection systems and other electronic systems. Google's Security Operations Center monitors Google's datacenter both locally and centrally. The GoGrid datacenter is protected by modern audio and video systems as well as local security guards. Microsoft combines an array of technologies to secure the physical integrity of its datacenter with cameras and alarms as well as by traditional lock and key means.

Amazon uses a two-factor authentication system to control employee access to its datacenters. Visitors must present identification and are permanently accompanied by authorized personnel throughout their visit. Everyone who goes in and out at Amazon is logged and regularly reviewed. Google only allows selected controlled and reviewed personnel access to its datacenters. Visitors are not allowed in Google datacenters at all. All GoGrid personnel must be registered and present valid ID before entering any of the company's buildings.

Amazon, GoGrid and Microsoft all use virtualization solutions in the cloud infrastructure. However, not all cloud vendors use the same solutions. Amazon and GoGrid, for example, use the Xen virtualization solution, while Amazon in contrast uses paravirtualization and GoGrid hardware virtualization. Microsoft, on the other hand, uses its own virtualization solution – Windows Azure Hypervisor.

Cloud vendors are confronted with numerous network attacks, such as distributed denial-of-service, man-in-the-middle or port scanning attacks, every day. A number of different applications and standard technologies are deployed to ward off these attacks. The defense mechanisms deployed by cloud vendors to fend off such attacks are described in brief in the following.

Amazon has its own methods of preventing successful distributed denial-of-service attacks which are not, however, described here. Microsoft prevents

[19] `http://www.gogrid.com/legal/sla.php`

denial-of-service attacks by using load balancing techniques to distribute workloads across several servers, and by deploying firewalls and intrusion prevention systems.

All Amazon APIs are available via SSL-protected endpoints which require server authentication. This prevents man-in-the-middle attacks in which the attacker attempts to obtain total control of data traffic and to inject and manipulate random information.

All inbound ports on Amazon EC2 instances are closed by default and therefore protected against port scanning. Any user can open any number of ports, however. Amazon stops and blocks port scanning as soon as it is detected.

Google is confronted by and attempts to disable the same kind of attacks by scanning its networks and applications with a number of different commercial and proprietary applications. Google also collaborates with third parties on the testing and improvement of the Google infrastructure and application security.

Most vendors use SSL and HTTPS to encrypt network connections. Access to Google Apps and most other Google end user programs is secured via an SSL connection. HTTPS access is also offered for most Google Apps services. Access to calendar and email can be set to HTTPS by default to restrict access to encrypted connections. Microsoft Office Live also offers an SSL connection which is not set as default but which can be activated at any time. GoGrid also offers SSL encrypted connections to its portal and for the API. Amazon Web Services can be reached via a secured SSL connection from the Internet and from within EC2.

6.2.2 Architecture

The data security strategies described in chapter 5.3 are now considered in connection with cloud vendors Amazon [5], Google [4] and FlexiScale[20]. The data encryption strategy is illustrated using the example of Amazon.

Amazon backs up Amazon S3, SimpleDB and Elastic Book Store data redundantly at several physical locations. The copies of Amazon Elastic Book Store are stored in the same Availability Zone and not across several zones. Google also backs up stored data redundantly across a large number of physical and logical storage capacities to ensure that data which has been unintentionally deleted can be recovered. FlexiScale also backs up data, but does not allow customers to recover virtual disks or individual files. Users must backup their own data themselves.

Data is not encrypted within Amazon Web Services. Data transfers can be encrypted, but the data is stored unencrypted. Service users can, however, encrypt data themselves before uploading it to servers and then store data in encrypted form.

[20] http://www.flexiscale.com/faqs.php

6.2.3 Administration

The service level agreements which govern use of a cloud service were described in chapter 5.4. All cloud vendors have their own service level agreements for the various services they offer. However, they are usually all very similar. Differences do exist in areas such as the issue of service credits for failure to meet service availability commitments. Some vendors offer service provisioning, i.e. they extend the contract term for the service by a specified number of days. Other vendors issue a credit which can be used to pay future fees. Examples of service level agreements are provided for the services Google Apps[21] and Amazon S3[22].

Google guarantees the availability of Google Apps for at least 99.9% of the time in any calendar month. If Google does not meet this obligation, the customer must request service credit within 30 days to avoid forfeiting the right to receive credit. Google's service credit cannot be converted to or exchanged for monetary accounts but entitles the customer to a maximum 15 days of added service. The downtime period is measured based on a server side error rate, whereby periods of less than 12 hours per calendar year are not counted towards downtime periods. The number of additional service days provided is shown in table 6.16. The monthly uptime percentage can be calculated – in simplified form – as follows:

$$\frac{\text{Total number of minutes in a calendar month - he number of minutes of downtime}}{\text{Total number of minutes in a calendar month}}.$$

Table 6.16: Google's service credit for failure to provide availability

Monthly uptime percentage	Days of service added
99,0% < x < 99,9%	3 days
95,0% < x < 99,0%	5 days
x < 95,0%	15 days

Like Google, Amazon S3 also guarantees availability for at least 99.9% of the time in any calendar month. If Amazon S3 does not meet this commitment, the user can request a service credit. If availability is between 99.0% and 99.9%, total charges are reduced by 10% and by 25% if availability is less than 99%.

The cloud vendors' service level agreements referred to also cover exclusions which are very similar in many ways. The following exclusion criteria in the SLAs of cloud vendors Amazon and Google are very similar. Downtime is not calculated if unavailability is caused by

1. Factors outside the control of either vendor,

[21] http://www.google.com/apps/intl/en/terms/sla.html
[22] http://aws.amazon.com/s3-sla/

2. Actions or inactions of the customer or any third party,
3. Customer equipment and/or third party equipment which is not within the direct control of the vendor.

6.2.4 Compliance

Chapter 5.5 has already discussed the need for data protection and privacy laws as well as security guidelines in cloud systems. The next section discusses cloud vendors which have security guideline certification and cloud vendors which take on board the data protection provisions under the Safe Harbor[23] Framework and the TRUSTe program[24].

Most cloud computing vendors are certified with the SAS (Statement on Auditing Standard) 70 Type II Report. This report must be produced for all outsourced services that impact company actions and confirms that an enterprise operates a functioning control system. Although SAS 70 is a U.S. standard, it is also important for many German and European enterprises which work for customers in the USA, for example. Amazon, Google, Microsoft, Salesforce and GoGrid are all SAS 70 certified cloud vendors.

In addition to SAS 70 Type II certification, Microsoft and Salesforce also hold an ISO/IEC 27001 certificate, the international standard for information security management systems (ISMS). This standard stipulates requirements for the implementation, monitoring, maintenance and improvement of a documented ISMS which may be certified to this standard. The certificate confirms that Microsoft and Salesforce have implemented the security mechanisms under this standard.

The Safe Harbor Framework governs data privacy principles which allow personal data to be transferred from the European Union to the United States of America. U.S. companies that register with the US Department of Commerce undertake to comply with certain European data protection requirements. Cloud vendors which have jointed the safe harbor system offer adequate protection in terms of notices, onward transfer, security, data integrity and access. Amazon, Google, Microsoft, IBM, Salesforce and Rackspace are examples of cloud vendors which have registered with the U.S. Department of Commerce.

The TRUSTe program exists alongside the Safe Harbor data privacy principles. TRUSTe is an independent, non-profit U.S. American initiative whose mission is to ensure that users must be asked for their permission before their data is used and are informed in a privacy statement about the following:

- What personal data is stored?
- How the data is used?

[23] `http://www.export.gov/safeharbor/`

[24] `www.truste.org`

- Whether data is passed on to third parties?
- What security measures are taken to prevent loss or misuse of, or changes to, data?
- How customers can inspect or change their data?

Microsoft, IBM and Salesforce are examples of cloud vendors who hold the TRUSTe seal.

Table 6.17 provides a complete list of the vendors discussed here and the certification they hold.

Table 6.17: Certificates of the vendors referred to

Vendor	TRUSTe	Safe Harbor	SAS 70 Type II	ISO/IEC 27001
Microsoft	x	x	x	x
Google	x		x	
Amazon	x		x	
Salesforce	x	x	x	x
PingIdentity			x	
Postini		x	x	
CohesiveFT				
Scalr				
RightScale				
IBM	x	x	x	x
GoGrid	x		x	
FlexiScale				
Rackspace	x			
LongJump				

The security technologies used in the domains of infrastructure, architecture, administration and compliance are not as yet sufficiently well documented to be able to check up precisely on the security measures adopted by cloud vendors. Information concerning administration is made available to cloud users in service level agreements, although these mainly detail service availability rather than security. Cloud vendors whose security guidelines have been certified, which have jointed the Safe Harbor system and/or the TRUSTe program publish this information on their websites. It is not always clear, however, how cloud vendors which do not hold these certificates or which have not joined programs proceed. A number of security technologies relevant to infrastructure have already been discussed, although there is still a lack of sufficient documented information in this area.

6.3 Application of the taxonomy to Amazon Cloud Services

The taxonomy in chapter 5 is now applied to the example of the cloud vendor Amazon. The answers to the checklists are based on public sources of information, such as websites, whitepapers [5] and the Amazon forum[25]. Questions are only dealt with here if they can be answered from the stated information sources. The application of the taxonomy is subdivided into the sections infrastructure, application and platform, administration and compliance.

6.3.1 Infrastructure

Physical security

Q: Do all the cloud vendor's datacenters operate to the same physical security standards?
A: All the cloud vendor's datacenters operate to the same physical security standards

Q: What physical security measures are taken by the datacenter which holds the cloud consumer's data?
A: The measures to guarantee the physical security of the datacenter taken by Amazon include: surveillance of sites and entrances to buildings by security guards, video camera surveillance, intruder detection systems and other electronic systems.

Q: In what way is the datacenter building itself secured?
A: The datacenter building is controlled by security guards, video camera surveillance and intruder detection systems.

Q: How are the entrances to the building secured?
A: Access to the building is controlled by a two-factor employee authentication system. Visitors must present identification and are continually accompanied by authorized personnel.

Q: What alarm systems are used?
A: Intruder detection systems are used.

Host

Q: What procedures are adopted to insulate the host?
A: The Xen hypervisor is used.

Q: Who has access rights to the hosts in the vendor's datacenter?
A: Only authenticated employees have access to the hosts in the datacenter.

Virtualization

[25] http://developer.amazonwebservices.com/connect/forumindex.jspa

Q: What virtualization technology does the cloud vendor use?
A: Amazon uses paravirtualization by Xen.

Network

Q: What procedures are adopted and network security systems used by a cloud vendor?
A: Amazon uses procedures such as SSL and firewalls.

Q: What technologies are used in order to stop network intrusions, such as denial-of-service attacks, man-in-the-middle attacks or port scanning?
A: Internal procedures, SSL and host based firewalls are all used to prevent attacks.

Q: How are these systems configured?
A: The firewall closes all ports by default.

Q: What configurations can or must the cloud consumer use?
A: The cloud consumer can configure the firewall ports.

6.3.2 Application and platform

Data security

Q: Where is data stored and how is it separated from other customers' data? Is the data on the cloud vendor's computers stored by the vendor in encrypted form?
A: The data is not stored by Amazon in encrypted form, but the cloud consumer can encrypt the data itself before it is stored.

Q: Where else is data stored, e. g. data backups and archiving or using a redundant cloud computing system?
A: Data is stored at several physical locations in the same region.

Q: If data is deleted, is it also deleted from all application instances, all caches and all backup copies?
A: All data and backup copies are deleted.

Q: What encryption procedures does the cloud vendor offer? Is the use of these procedures stipulated by contract?
A: Amazon does not use any encryption procedures, but the cloud consumer can encrypt and store data itself.

Q: Can backup copies be encrypted?
A: No, these are only encrypted if the data itself is already encrypted.

Q: Can data be recovered after it has been deleted?
A: Data cannot be recovered as it is securely erased.

Q: Can stored data be shared with other cloud consumers?

A: The data can be enabled for defined Amazon users or for general use.

Application security

Q: What authentication mechanisms are offered and are these mechanisms appropriate to the sensitivity of the data?
A: Amazon offers multi-factor authentication.

Security as a service

Q: Is the security architecture documented in full?
A: No, the security architecture is not fully documented.

6.3.3 Administration

Phases of service use

Q: Are all the vendor's security functions documented? Is enough information available for assessment purposes?
A: Not enough information is available for assessment purposes given that not all security functions are documented.

Q: What guarantees does the vendor's standardized service level agreement provide? What exceptions are there?
A: Amazon guarantees 99.9% availability in any calendar month. Exclusions apply to downtimes, for example, which do not count towards service credits if the problems are due to factors outside the control of Amazon, are due to actions or inactions of the customer or any third party, or due to customer equipment and/or third party equipment.

Q: Where and with which system components is the performance of a cloud service monitored? Is it possible to integrate third party services?
A: Third party services, such as VPN Cubed for EC2, can be integrated in Amazon.

Identity and rights management

Q: What rights issue and controlled rights withdrawal processes are used?
A: Rights are issued using Access Control Lists (ACL).

Q: Is there a programming interface for the provision and deletion of rights?
A: ACLs can be modified via the API.

6.3.4 Compliance

Data privacy

Q: What sort of data privacy rules does the cloud service apply? Can the vendor provide a copy of the guidelines in document form?
A: Amazon holds the TRUSTe seal.

Q: Who has access to the cloud computing system host?
A: Access to the cloud computing system hosts is restricted to those administrators who absolutely require access for operational purposes.

The legal framework

Q: Can the geographical location be restricted to the extent necessary to meet statutory requirements?
A: The geographical locations at which data is stored can be restricted.

Governance

Q: What certificates does the cloud vendor hold? What security aspects are covered by these certificates?
A: Amazon holds the SAS 70 Type II certificate.

Q: Are certificates only audited internally, or are they audited externally as well?
A: Audits are performed by external bodies.

6.4 Conclusion

It is very difficult to assess the implementation of security measures by cloud vendors given the paucity of information provided by the vendors themselves. Many vendors document the use of SSL and HTTPS but do not offer information about any other technologies used. Standard cloud system technologies should be defined and introduced in the near future.

Cloud vendors usually provide information about service level agreements and the safeguarding of privacy. However, this information is often so vague that it is only possible to speculate on the way in which it is actually used. When it comes to assessing user privacy a distinction must be made between those enterprises which store data in the European Union and those which store data in the USA – different rules and regulations apply in each case. The Safe Harbor Framework and the TRUSTe program are helpful in this context, although not all cloud vendors – i.e. CohesiveFT or RightScale – have accepted these data privacy principles. Another important point is the measurement of time and volume based values and the monitoring of contractually agreed service quality. Cloud vendors already use metering procedures, but the measured values are not transparent for cloud users.

Until such time as statutorily mandatory standards apply to cloud systems, users would be well advised to assess all cloud vendors very carefully before using their services. Such an assessment should cover the most important security aspects discussed in chapter 5.

7 Summary and outlook

In recent years cloud computing has become an important buzzword referring to the provisioning of IT services on remote resources and their procurement from, in most cases, public networks. Cloud proliferation is accompanied by a permanent revolution in the services launched on the market by vendors. In this context particular attention is merited by the security functions which are offered by cloud services.

The discussion in the previous chapter reveals a very mixed picture as far as the security aspects of cloud services are concerned. Essential security functions which use known technologies are also used in cloud computing systems to encrypt a data channel, for example. Vendors do differ quite radically in some cases, however, in terms of the security features they support. The lack of a standardized security configuration also makes it difficult to compare different vendors.

The following sections summarize the findings of the study and look ahead to some of the issues which will need to be resolved in order to provide efficient and user-friendly cloud computing services in the future. This section concludes by outlining the services offered by Fraunhofer AISEC in the field of cloud computing security.

7.1 Study findings

The study comes to the following conclusions:

- The structure of cloud computing systems comprises four layers – end user, software, platform, and infrastructure – and the players acting on these layers form a very complex IT security framework. This study describes all the key layers and players that must be examined, depending on the application and the selected cloud service.
- Certified tools based on cloud services are essential for cloud computing systems to increase the portability and interoperability of individual cloud service offerings. Standardization bodies, reference implementations, and development environments adapted to cloud computing systems must exist for this purpose.

- The cloud security taxonomy provides a clearly structured framework of the security areas that should be considered when using cloud services. Owing to the rapid development pace of both the technologies and the existing services, the application of the cloud taxonomy should be project based and the weighting accorded to individual security areas adapted to the specific requirements in each case.

- Modern cloud service portfolios clearly use a whole series of security technologies already, especially on the infrastructure layer. On the other hand, when it comes to architecture, administration, and compliance, cloud vendor support for security technologies is not yet adequate to achieve the stipulated protection goals. More detailed analyses are called for here to identify which current technologies are potentially suitable and determine whether new technologies need to be developed. There is a trend toward procuring certain security functions, such as parts of the identity or access management functionality, as a service from specialist vendors.

- On the administration side, service level agreements are an important instrument for specifying all the rights and obligations that exist between cloud users and cloud vendors. The standardized service level agreements offered at present, which are not normally freely negotiable by cloud users but can simply be either accepted or rejected, provide only minimal guarantees regarding cloud service quality. In particular, the security guarantees contained in these agreements are very rudimentary, and need to be extended in order to achieve the above-mentioned protection goals. Systems to facilitate automatic monitoring and testing of the agreed service quality criteria are also essential.

- From the compliance perspective, there are no objections to the use of cloud services. However, the responsibility for the data concerned usually lies with the cloud user, who needs to define precise guidelines stating which information is allowed to be stored and processed in a cloud service and how, and simultaneously specifying the necessary security functions. From a legal viewpoint, too, the restrictions to which certain data is subject and the use of specific cloud services should be separately considered in each case.

- The market overview in chapter 6.1 gives a general rundown of selected cloud services together with their prices and functionalities. The taxonomy of secure cloud computing is then applied to these services and their security functions assessed. It is fair to say here that information about the implemented security functions is not adequately documented by cloud vendors. In many cases, security plays only a minor role when they present their services, so that detailed information should be requested from the vendor upfront of choosing or using a specific cloud service. If appropriate, a proof of concept should be realized before the service is actually put to productive use.

7.2 Unresolved issues

Unresolved issues mainly concern the architecture, administration and taxonomy compliance of secure cloud computing, on which – as already briefly discussed in the conclusions to the study presented above – a detailed analysis still needs to be performed. Many of the existing cloud services on offer have already gone into standard operation in a number of different areas. However, the question arises as to which criteria should be applied when choosing and evaluating services. Integration in existing systems is another issue which has not been entirely resolved.

Continuous security checks of the kind already performed in companies themselves should also be carried out on cloud services. Guidelines and standardized procedures need to be defined for this purpose to ensure that checks can be undertaken efficiently. Various security levels could, for example, be stipulated, wholly independently of the cloud user's data and processes, and subject to separate audits.

Fraunhofer AISEC is responding to the increasing sophistication of security technologies for cloud services by setting up a test environment for cloud services in which various security configurations can be run, both within the Fraunhofer AISEC's own cloud computing system and in public cloud computing systems.

7.3 Fraunhofer AISEC services

The Fraunhofer AISEC develops tailored, directly applicable solutions for all branches of business and industry, including the health, transport, traffic, logistics and production sectors as well as for commerce and financial services. In addition to client oriented contract research, Fraunhofer AISEC offers consultation on the development of security concepts and conducts studies for both the private sector and government.

Fraunhofer AISEC's cloud computing activities are managed by Dr. Werner Streitberger, and are concentrated in the "Secure Services and Quality Testing" department at its new site in Garching near Munich. The department's work includes providing consulting on the use of cloud services, implementing proof of concepts and the cloud test laboratory.

- Concept development and consulting for the use of cloud services: In the field of concept development and consulting for the use of cloud services the department assesses current cloud services according to financial and technical criteria and selected to meet specific requirements. If necessary, services are compared with non-cloud based solutions and with existing systems and services. The focus here is on the security of the solution and the choice of technologies which can be presented in the framework of workshops and studies.

- Implementation of proof of concept solution: As well as developing architecture concepts, the security aspects of integration in existing processes are examined and implemented in the form of prototypes. Customers also receive support starting up their cloud based solutions.
- Cloud test laboratory: Drawing on the many years of experience of its laboratory personnel, the still evolving Fraunhofer AISEC cloud test laboratory offers tests of existing cloud service installations and issues security certificates. Private and hybrid cloud configurations can also be developed and security checks performed on the laboratory's own hardware. Demonstrators and prototypes can also be developed as part of the cloud test laboratory's activities with the aim of performing compliance and interoperability tests, for example.

Bibliography

[1] *Information technology - open system interconnection - basic reference model: The basic model*, 1994. 43

[2] *Web services security*, March 2004. 45

[3] *Xml schema part 0: Primer second edition*, October 2004. 45

[4] *Comprehensive review of security and vulnerability protections for google apps*, February 2007. 78, 79

[5] *Amazon web services: Overview of security processes*, June 2009. 78, 79, 83

[6] *Securing microsoft's cloud infrastructure*, May 2009. 78

[7] Almond, Carl: *A practical guide to cloud computing security - what you need to know now about your business and cloud security*. Technical report, Avanade Inc., August 2009. 40

[8] Amrhein, Dustin, Andrew de Andrade, Joe Armstrong, Ezhil Arasan B, Richard Bruklis, Ken Cameron, Reuven Cohen, Rodrigo Flores, Gaston Fourcade, Thomas Freund, Babak Hosseinzadeh, William Jay Huie, Sam Johnston, Ravi Kulkarni, Anil Kunjunny, Gary Mazzaferro, Andres Monroy-Hernandez, Dirk Nicol, Lisa Noon, Santosh Padhy, Thomas Plunkett, Ling Qian, Balu Ramachandran, Jason Reed, German Retana, Dave Russell, Krishna Sankar, Alfonso Olias Sanz, Wil Sinclair, Erik Sliman, Patrick Stingley, Robert Syputa, Doug Tidwell, Kris Walker, Kurt Williams, John M Willis, Yutaka Sasaki, Eric Windisch, and Fred Zappert: *Cloud computing use cases*. Technical report, Cloud Computing Use Case Discussion Group, July 2009. 11, 13, 25

[9] Armbrust, Michael, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy H. Katz, Andrew Konwinski, Gunho Lee, David A. Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia: *Above the clouds: A berkeley view of cloud computing*. Technical Report UCB/EECS-2009-28, Electrical Engineering and Computer Sciences, University of California at Berkeley, February 2009. 26, 34, 41, 59

[10] Bardin, Jeff, Jon Callas, Shawn Chaput, Pam Fusco, Francoise Gilbert, Christofer Hoff, , Dennis Hurst, Subra Kumaraswamy, Liam Lynch, Scott Matsumoto, Brian Higgins, Jean Pawluk, George Reese, Jeff Reich, Jeffrey Ritter, Jeff Spivey, and John Viega: *Security guidance for critical areas of*

focus in cloud computing. Technical report, Cloud Security Alliance, April 2009. 37, 39, 49

[11] Bernstein, David, Erik Ludvigson, Krishna Sankar, Steve Diamond, and Monique Morrow: *Blueprint for the intercloud - protocols and formats for cloud computing interoperability*. Internet and Web Applications and Services, International Conference on, 0:328–336, 2009. 49

[12] Borthakur, Dhruba: *The Hadoop Distributed File System: Architecture and Design*. The Apache Software Foundation, 2007. 33

[13] Cavoukian, Ann: *Privacy in the clouds*. Technical report, Information and Privacy Commissioner of Ontario, 2009. 58

[14] Chakrabarti, Anirban: *Grid Computing Security*. Springer, Berlin, Juni 2007. 41

[15] Comission, European: *Directive 95/46/ec of the european parliament and of the council of 24 october 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*. Official Journal of the European Communities, 281:31, 1995. 57

[16] Dean, Jeffrey and Sanjay Ghemawat: *Mapreduce: Simplified data processing on large clusters*. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation*, pages 137–150, San Francisco, CA, December 2004. `http://www.usenix.org/events/osdi04/tech/dean.html`. 33

[17] Eckert, Claudia: *IT-Sicherheit*. Oldenbourg, 6. edition, 2009. 18, 19, 21

[18] Erdogmus, Hakan: *Cloud computing: Does nirvana hide behind the nebula?* IEEE Software, 26(2):4–6, 2009. 4

[19] Fink, Simon: *Datenschutz zwischen Staat und Markt : die Safe-Harbor-Loesung als Resultat einer strategischen Interaktion zwischen der EU, den USA und der IT-Industrie*. PhD thesis, Uni Konstanz, 2003. 58

[20] Gellman, Robert: *Privacy in the clouds: Risks to privacy and confidentiality from cloud computing*. Technical report, World Privacy Forum, February 2009. 19

[21] Greenberg, Albert, James Hamilton, David A. Maltz, and Parveen Patel: *The cost of a cloud: research problems in data center networks*. SIGCOMM Comput. Commun. Rev., 39(1):68–73, 2009, ISSN 0146-4833. 39

[22] Grossman, Robert L.: *The case for cloud computing*. IT Professional, 11(2):23–27, 2009, ISSN 1520-9202. 6

[23] Hayes, Brian: *Cloud computing*. Communications of the ACM, 51(7):9–11, 2008, ISSN 0001-0782. 4

[24] Heiser, Jay and Mark Nicolett: *Assessing the security risks of cloud computing*. Technical Report G00157782, Gartner Research, June 2008. 37

[25] Horrigan, John B.: *Data memo*. Technical report, PEW Internet and American Life Project, September 2008. 28

[26] Leavitt, Neal: *Is cloud computing really ready for prime time?* Computer, 42(1):15–20, January 2009. 4, 8

[27] Lin, Geng, David Fu, Jinzy Zhu, and Glenn Dasmalchi: *Cloud computing: It as a service*. IT Professional, 11(2):10–13, 2009, ISSN 1520-9202. 4

[28] Mell, Peter and Tim Grance: *Darft nist working definition of cloud computing*. Technical Report Version 15, National Institute of Standards and Technology, Information Technology Laboratory, August 2009. 4, 5, 24

[29] Mowbray, Miranda: *The fog over the grimpen mire: Cloud computing and the law*. Technical Report HPL-2009-99, HP Laboratories, 2009. 49

[30] Pearson, Siani and Andrew Charlesworth: *Accountability as a way forward for privacy protection in the cloud*. Technical Report HPL-2009-178, HP Laboratories, 2009. 23

[31] Pfitzmann, Birgit and Michael Waidner: *Security Protocols*, volume Volume 3364/2005 of *Lecture Notes in Computer Science*, chapter Federated Identity-Management Protocols, pages 153–174. Springer Berlin / Heidelberg, 2004. 52

[32] Recordon, David and Drummond Reed: *Openid 2.0: a platform for user-centric identity management*. In *DIM '06: Proceedings of the second ACM workshop on Digital identity management*, pages 11–16, New York, NY, USA, 2006. ACM, ISBN 1-59593-547-9. 52

[33] Ristenpart, Thomas, Eran Tromer, Hovav Shacham, and Stefan Savage: *Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds*. In *Proceedings of CCS 2009*. ACM Press, November 2009. 47

[34] Smith, Matthew: *Security for Service-Oriented On-Demand Grid Computing*. PhD thesis, Fachbereich Mathematik und Informatik, Universität Marburg, 2008. 34

[35] Staten, James, Simon Yates, Frank Gillett, Walid Saleh, and Rachel A. Dines: *Is cloud computing ready for the enterprise?* Technical report, Forrester Research, Inc., March 2008. 8

[36] Stock, Andrew van der, Jeff Williams, and Dave Wichers: *Owasp top 10: The ten most critical web application security vulnerabilities*. Technical report, OWASP Foundation, 2007. 46

[37] Streitberger, Werner: *Einsatz von Risikomanagement bei der Steuerung von Grid-Systemen - Eine Analyse von Versicherungen anhand einer simulierten Grid-Ökonomie*. PhD thesis, Lehrstuhl für Wirtschaftsinformatik, Fakultät für Rechts- und Wirtschaftswissenschaften, Universität Bayreuth, 2009. 8, 28

[38] Varia, Jinesh: *Cloud architectures*. Technical report, Amazon Web Services, 2008. 28

[39] Vishwanath, Kashi Venkatesh, Albert Greenberg, and Daniel A. Reed: *Modular data centers: how to design them?* In *LSAP '09: Proceedings of the 1st ACM workshop on Large-Scale system and application performance*, pages 3–10, New York, NY, USA, 2009. ACM, ISBN 978-1-60558-592-5. 39

[40] Wisniewski, Thomas, Tony Nadalin, Scott Cantor, Jeff Hodges, and Prateek Mishra: *Saml v2.0 executive overview*, April 2005. 52

[41] Zaharia, Matei, Dhruba Borthakur, Joydeep Sen Sarma, Khaled Elmeleegy, Scott Shenker, and Ion Stoica: *Job scheduling for multi-user mapreduce clusters*. Technical Report UCB/EECS-2009-55, Electrical Engineering and Computer Sciences, University of California at Berkeley, April 2009. 41

Contact Details

Fraunhofer AISEC
(formerly Fraunhofer SIT-Munich)
Parkring 4
D-85748 Garching b. München

Tel.: +49 (0)89-322-9986-0
Fax: +49 (0)89-322-9986-299
http://www.aisec.fraunhofer.de
http://www.cloudsecuritylab.de

Planning and realisation of the study

Dr. Werner Streitberger (Lead)

Angelika Ruppel
Tel.: +49 (0)89-322-9986-154
angelika.ruppel@aisec.fraunhofer.de

Area of research "Secure Services and Quality Testing"

Mario Hoffmann
Tel.: +49 (0)89 322-9986-177
mario.hoffmann@aisec.fraunhofer.de

Press and Public Relations

Oliver Küch
Tel.: +49 (0)6151869213
oliver.kuech@sit.fraunhofer.de

State of study: September 2009