

SECURITY IN SERVICE LEVEL AGREEMENTS FOR CLOUD COMPUTING

Karin Bernsmed, Martin Gilje Jaatun

SINTEF Information and Communication Technology, Trondheim, Norway

Karin.Bernsmed@sintef.no, Martin.G.Jaatun@sintef.no

Astrid Undheim

Telenor Corporate Development, Trondheim, Norway

Astrid.Undheim@telenor.com

Keywords: Cloud computing, security, Service Level Agreements, SLA, Quality of Service, QoS.

Abstract: The Cloud computing paradigm promises reliable services, accessible from anywhere in the world, in an on-demand manner. Insufficient security has been identified as a major obstacle to adopting Cloud services. To deal with the risks associated with outsourcing data and applications to the Cloud, new methods for security assurance are urgently needed. This paper presents a framework for security in Service Level Agreements for Cloud computing. The purpose is twofold; to help potential Cloud customers to identify necessary protection mechanisms and, in the next step, to facilitate automatic service composition based on a set of predefined security requirements. We demonstrate the practical applicability of the first objective with a small case study.

1 INTRODUCTION

Cloud computing is a hot topic in both industry and research. Cloud services can be found everywhere, offering all possible IT services imaginable in an on-demand and scalable manner. In the Cloud, security has been identified as one of the main obstacles for adoption. This contrasts with the fact that few of the security issues related to cloud computing are new and unique; most of them have been investigated and addressed in the traditional system and network security context. Mechanisms for data protection, access control, trust delegation and mitigation of DDoS attacks, etc. are well known and have been (more or less) successfully applied in many other contexts. However, some characteristics of cloud computing are fundamentally new in an outsourcing perspective, such as multi-tenancy and on-demand elasticity, which creates new threats.

Cloud computing is outsourcing, and outsourcing implies trust. In the Cloud the responsibility for implementing and maintaining efficient security mechanisms will be in the hands of the provider. To overcome the fear of the Cloud, the provider needs to convince the customer that his data and applications will be properly secured. Trusting the provider is one thing; however, in the near future it is expected

that complex Cloud services will be composed of several other services from different providers in multiple steps, something that will introduce a longer chain of trust than the traditional ICT system outsourcing context. It will be a huge challenge to assure the customer that adequate security mechanisms exist and are correctly implemented throughout the whole chain of providers.

To mitigate the security risks associated with the Cloud, existing security mechanisms and their effectiveness should be formalized in contracts. In this paper we argue that Service Level Agreements (SLAs) in the Cloud should be extended to include the security mechanisms offered by the provider. We then outline a framework for such mechanisms. The purpose is not only to increase the trust in the provider, but also facilitate for potential Cloud customers to make an objective comparison between different service providers on the basis of their security features. Moreover, identifying and organizing critical security mechanisms in a systematic manner forms a basis for composing services from different providers, based on a set of predefined security requirements.

This paper is organized as follows. We discuss some of the previous work related to security management in the Cloud in Section 2. Section 3 explains our vision of SLAs for Cloud computing. In Sec-

tion 4 we identify security mechanisms suitable for Cloud services and outline a framework for security-aware Cloud SLAs. Section 5 demonstrates how a real-world scenario could benefit from our approach. Section 6 concludes the paper by discussing our approach and laying out the directions for future work.

2 RELATED WORK

Security mechanisms for Cloud computing has previously been discussed in a comprehensive report published by ENISA (European Network and Information Security Agency (ENISA), 2009), which outlines an information assurance framework where division of responsibilities between the customer and the Cloud provider is clarified. Another example is the report by Gartner (Heiser and Nicolett, 2008) that provides a basic description of the potential risks involved with Cloud computing. We have used these reports as a basis for the security framework derived in this paper.

There have been some projects in the research community looking into various aspects of security in SLAs. Early work on security agreements was performed already in 1999 by Henning (Henning, 2000), who already then raised the question whether security can be adequately expressed in a SLA. Security requirements for web services have been treated by Casola et. al., who proposed a methodology to help evaluate and compare security SLAs (Casola et al., 2006). Frankova and Yautsiukhin have also recognized the need for security in SLAs (Frankova and Yautsiukhin, 2007), however, their approach focuses on the process of selecting the optimal service composition based on a set of predefined requirements rather than detailing what security mechanisms to include. Also de Chaves et al. explore security in SLAs, but with focus on measurable security metrics (de Chaves et al., 2010), combined with a monitoring and controlling architecture (Righi et al., 2006). Our approach aims to be more practically applicable, since we provide a framework for auditable security mechanisms specifically targeted for Cloud computing services SLAs.

3 SERVICE LEVEL AGREEMENTS FOR CLOUD COMPUTING

Deploying services in the Cloud creates new challenges for both service providers and customers, especially regarding the service quality. The customers

have less control of the service delivery, and need to take precautions in order not to suffer low performance, long downtimes or loss of critical data. Service Level Agreements (SLAs) have therefore become an important part of the Cloud service delivery model. A SLA is a binding agreement between the service provider and the service customer, which is used to specify the level of service to be delivered. A SLA includes information about the service delivered by the cloud provider, together with the penalties if the SLA is broken. The penalties are usually stated as service credits to the customers.

3.1 Quality of Service in the Cloud

In order to verify that the Cloud provider delivers service in accordance with the SLA, the SLA usually contains Quality of Service (QoS) parameters. QoS refers to the (measurable) ability of a distributed system to provide the network and computation services such that the customer's expectations are met. SLAs have previously been used for many years to specify QoS guarantees between telecom operators and corporate customers, as well as for regulating outsourcing contracts. However, SLAs have (until now) not been adopted by the public at large.

The advent of the Cloud concept, which is characterized by its elastic and measurable services (Mell and Grance, 2009), has given rise to a new demand for such agreements. In the Cloud, two aspects of QoS are of special interest; *dependability* and *performance*. Service dependability is usually defined as a combination of the service availability (the proportion of time a system delivers service according to the requirements) and reliability (the ability to provide uninterrupted service) (International Telecommunication Union, 2008), whereas performance is usually characterized by throughput (the number of bits per second of data transmitted or processed) and delay (the number of seconds used for transmission or processing). The term QoS usually does not include security, even though some previous efforts have tried to extend the term in this respect (Irvine, 2000; Lindskog, 2005).

Today, most of the major Cloud service providers include QoS guarantees in their SLA proposals, however the focus in most cases is on availability¹. In most cases, the SLA lacks performance guarantees, which from the customer's point of view is a major drawback. A very low performance will be perceived

¹An example is the Amazon EC2 Cloud service, which at the time of writing offers 99.95% availability on a yearly basis and issues 10% credits if the SLA is broken. Performance is not mentioned.

by the customer as service unavailability and should be credited accordingly.

3.2 Security in SLAs

Even though service availability and performance often are identified as critical issues, the number one barrier of adopting Cloud computing services is assurance (European Network and Information Security Agency (ENISA), 2009): how can a potential customer be sure that it is safe to place data and applications in the Cloud? Since the SLA is used to explicitly state the obligations of the provider, the implemented security mechanisms, their effectiveness and the implications of possible mismanagement should be a part of this agreement. This concept is also known as Quality of Protection (QoP), which comprises the ability of a service provider to deliver service according to a set of specific security requirements. A standardized framework for constructing a SLA in the Cloud, based on guaranteed levels of these attributes and the consequences of mismanagement, is therefore of utmost importance for creating trustworthy and reliable Cloud computing services. This includes clarifying the consequences of a service provider’s possible failure to deliver the service in accordance with the contract. Figure 1 outlines the basic structure for such a SLA.

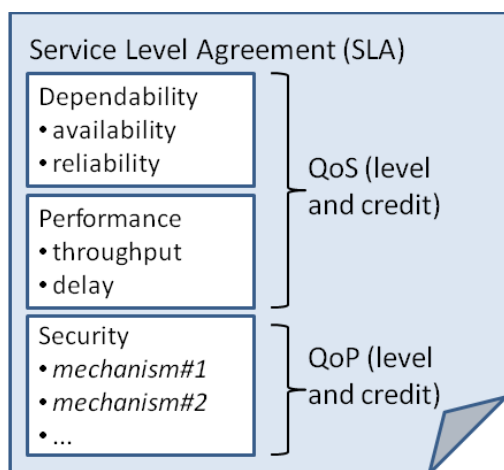


Figure 1: The basic structure of a SLA with dependability, performance and security guarantees

The specific security mechanisms that should be included in the SLA will then mirror the needs for the particular (composed) service.

3.3 Auditing Security Effectiveness

A plausible objection against introducing security in a SLA (as proposed in Figure 1) derives from the inherent difficulty in *measuring* security (Verendel, 2009); it is generally very difficult to prove that, e.g. an encryption algorithm is secure, but often very easy to show that it is insecure when a flaw has been discovered. Another option is therefore to measure breaches that occur, rather than trying to measure security strength. This means that our framework do not require the provider to estimate the strength of their security mechanisms, but simply expects them to provide the agreed level of security in any way they see fit. If they fail, the punitive mechanisms come into play.

In practice, it is often up to the Cloud customer to monitor and verify that the delivered service is in accordance to the QoS guarantees stated in the contract. This means that transparent monitoring solutions are needed in order for the customer to measure or check the provided service level. Methods for auditing the security mechanisms and measuring breach frequencies will be necessary to use our framework. We discuss this issue further in Section 6.

4 SECURITY MECHANISMS FOR CLOUD SLAs

Returning to the Cloud SLA outlined in Figure 1, in order to identify what security mechanisms to include in the contract we need to take a further look into the security related threats associated with the Cloud. As explained in the introduction, Cloud computing presents some fundamentally new security challenges in addition to the traditional ones. Following the well established definition from NIST (Mell and Grance, 2009), Cloud computing has five important characteristics, namely: i) on-demand self-service, ii) broad network access, iii) resource pooling, iv) rapid elasticity, and v) measured service. The on-demand aspect must be taken into account when providing Cloud SLAs, since a SLA must also be allowed to be composed on-demand. Our framework is specifically constructed with this in mind. Network access does not present any new security challenges but network security becomes even more important in the Cloud, where large amounts of confidential data are regularly transmitted over the public Internet. Resource pooling and rapid elasticity presents some new challenges with respect to security, as will be described in the following. In addition, we have identified three categories of security mechanisms that need

special attention in Cloud computing, and therefore should be a part of the SLA. These are "access control", "audit, verification and compliance" and "incident management and response". Together with "secure resource pooling" and "secure elasticity" these five categories can be used in a structured approach to pick the right security mechanisms for a particular service.

Figure 2 illustrates our proposed framework. The security mechanisms are divided into three main service categories, depending on the particular Cloud resources that are used. These are storage, processing (CPU and memory), and network. These three resources can be virtualized and offered as Infrastructure as a Service (IaaS), they can be used to offer a Platform as a Service (PaaS), or applications can run on top of these physical resources and being offered as Software as a Service (SaaS). The security mechanisms suggested for the framework are related to Secure Resource Pooling (RP), Secure Elasticity (E), Access Control (AC), Audit, Verification & Compliance (AU) and Incident Management & Response (IM); these will be further described below.

4.1 Secure Resource Pooling

Resource pooling in Cloud computing today is achieved by using virtualization either at the hardware level (as with Amazon Web Services²) or at the application level (as with Google Docs³). Both techniques enable multi-tenancy, i.e., different users share the same resources, and virtualization ensures the isolation of data and applications owned by different users.

The sharing of physical resources in the Cloud give rise to new security threats. One of the most imminent is unauthorized access to applications or data through the hypervisor (Christodorescu et al., 2009), which may occur if proper isolation of applications and data is not achieved. It is therefore necessary to make sure that protection mechanisms exist and that they are stated in the SLA. In the framework outlined in Figure 2 this is illustrated as "RP1: Data isolation" and "RP8: Application isolation", which are related to storage services and processing services, respectively. Moreover, resource sharing implies that the customers need guarantees that their property remains confidential (RP3: Data encryption) and is integrity protected (RP6: Data integrity, RP12: Application integrity), that their data and applications are properly deleted from the physical hardware when requested (RP2: Data deletion), and that the data can be brought back in-house if necessary (RP5: Data portability,

²<http://aws.amazon.com>

³<http://docs.google.com>

RP7: Data back-up). The customer should also have the possibility to put restrictions on the geographic location of storage and processing (RP4: Data location, RP9: Application location). Regarding network services (inside Clouds, between Cloud data centers and between the Cloud and the customer's premises), the customer should make sure that his traffic is properly protected (RP13: Network encryption, RP15: Integrity protection) and isolated from other customers traffic (RP14: Traffic isolation).

4.2 Secure Elasticity

Cloud computing promises rapid scalability of resources, scaling up and releasing resources as needed. This elasticity is also enabled by virtualization. Adding more virtual resources on the same physical machine does not in itself pose any new threats, but migrating virtual resources to new physical resources requires a secure migration process (E1: Secure data migration, E2: Secure virtual machine migration), including the actual network transfer. It must also be ensured that the new physical resource fulfils the same security requirements.

4.3 Access Control

Access control is especially important in the Cloud, where both competing customers sharing the same resources as well as insider personnel may try to gain unauthorized access to the customer's data. The resource must also be protected from unauthorized remote access. It is therefore crucial to make sure that proper access control mechanisms are implemented (AC1: Identity management, AC2: Access management, AC3: Key management), and that there are strict restrictions on e.g. who may enter Cloud datacenters (AC4: Internal security control).

4.4 Audit and Verification

The possibility to audit and verify the security of a service is very often crucial to the customer, however in the Cloud this is often not standard practice. Customers may require access to server logs, failed login attempts records or database change records (AU1: Logging) and sometimes also the possibility to audit the activity on specific Cloud resources (AU2: Auditing). In addition, the customers may want to make sure that a security certification scheme exists and is adapted to the Cloud infrastructure (AU3: Certification). Customers may also have privacy concerns (AU4: Customer privacy).

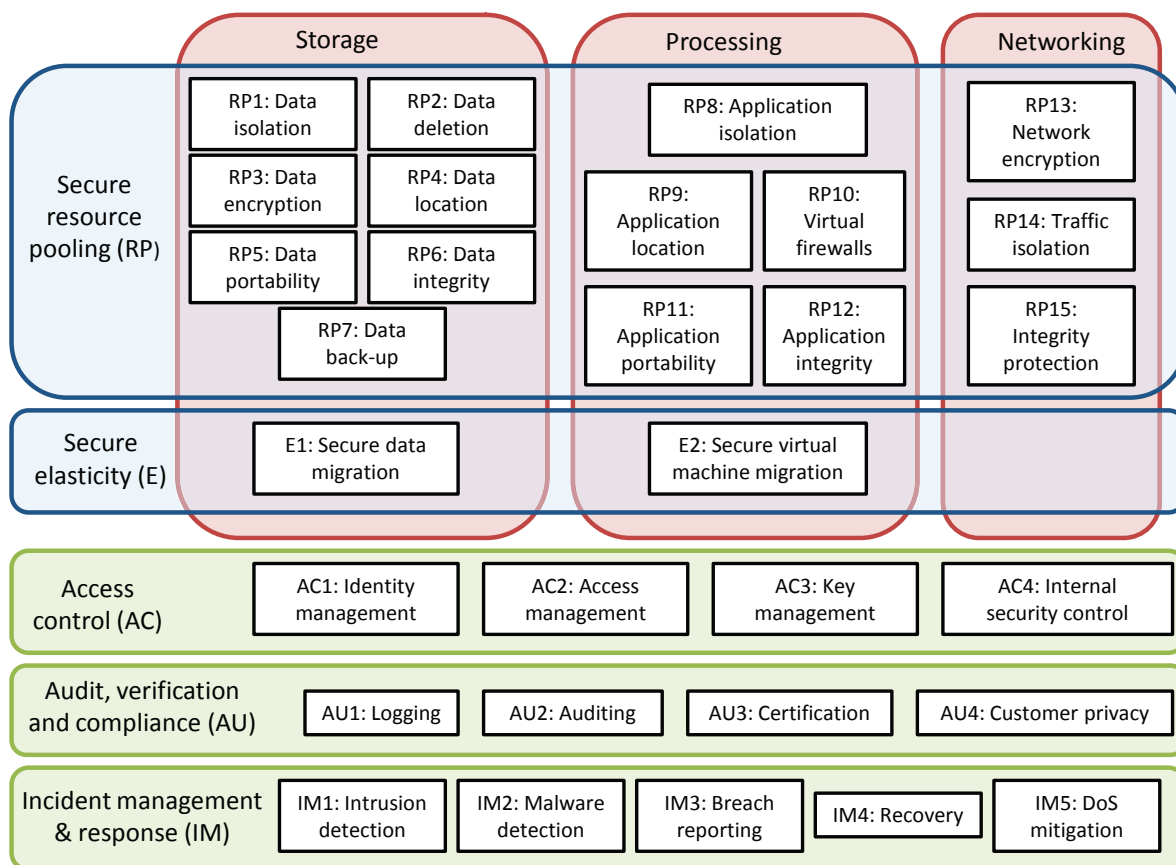


Figure 2: A framework for security mechanisms for Cloud SLAs

4.5 Incident Management and Response

To make sure that the Cloud provider detects and responds to threats, the SLA may contain mechanisms for intrusion and malware detection (IM1: Intrusion detection, IM2: Malware detection), that security breaches are recorded and reported (IM3: Breach reporting), that data and applications can be reconstructed in the case of disasters (IM4: Recovery) and that mechanisms to prevent and mitigate DoS attacks are implemented (IM5: DoS mitigation).

The framework in Figure 2 represents a first step towards a security-aware SLA, as we described it in Section 3. The purpose of the framework is to serve as a basis for constructing a SLA for a specific Cloud service, by identifying security mechanisms that should be stated in the SLA. In the next section we will demonstrate how the framework can be applied in a real-world scenario.

5 CASE STUDY: CUSTOMER RELATIONSHIP MANAGEMENT

Customer Relationship Management (CRM) represents the total set of processes and applications for interacting with customers, both for sale, customer service, and marketing. When applying Payne & Frow’s preferred perspective of CRM as “a holistic approach to managing customer relationships to create shareholder value” (Payne and Frow, 2005), the software solution is often quite complex. Maintaining this software is typically not part of “core business” of the CRM user, and CRM applications is thus one area where companies have successfully embraced the SaaS concept.

Payne & Frow presents a framework for CRM strategy composed of five main processes:

- Strategy Development
- Value Creation
- Multichannel Integration

- Performance Assessment
- Information Management

In principle, a CRM system can cover the complete sale process for a business, implying that the customer-related information stored in the system is absolutely vital to the business, and exposure of this information to a competitor would be a serious threat. In addition, private customer data such as address and account information should not be visible to others.

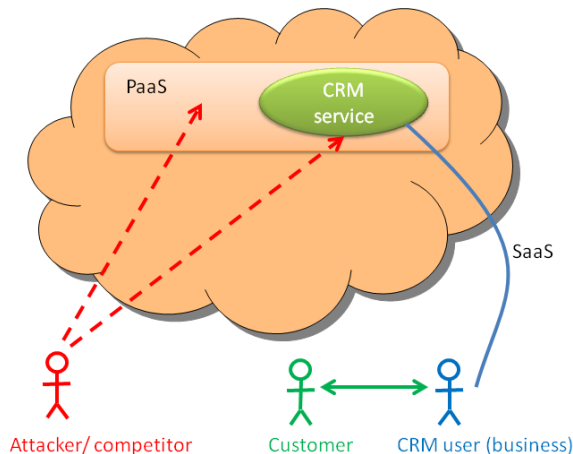


Figure 3: A Customer Relationship Management case

The Cloud-based CRM system is illustrated in Figure 3. We assume that the CRM user (the business) interacts with its customers via some undetermined channel, but that the customers are not generally aware of the CRM application being Cloud-based (in general, the customers may not even be aware of there being a CRM application at all). The CRM user is likely to have a competitor, and it is quite possible that the competitor is a user of the same CRM service – or could be a user of an underlying PaaS offered by the same provider.

The multi-tenancy property of Cloud services implies that it is necessary to protect against both attacks that are completely external, and against attacks from other users of the service. As mentioned, the latter attackers may have legitimate access to the Cloud provider’s infrastructure either at SaaS or PaaS levels.

Next, the elasticity property of Cloud computing means that both applications and stored data are (more or less) frequently migrated between physical resources. This is true also for CRM data and applications. A company may typically see peaks in customer interaction around holidays or after some disaster (for a travel agency, insurance company etc.).

Referring to Figure 2, there are a number of security mechanisms that are important to the CRM user when deploying the CRM application in a Cloud. The

most important mechanisms are related to secure resource pooling, most importantly for the stored customer data:

- Data isolation (RP1) - stored customer data must not be available to competitors.
- Data deletion (RP2) - data no longer needed must be properly deleted.
- Data encryption (RP3) - stored customer data must be encrypted.
- Data location (RP4) - parts of the customer data may be defined as Personal Identifiable Information (PII) and must be treated thereafter.
- Data portability (RP5) - the CRM user wants to avoid lock-in of customer data.
- Data integrity (RP6) - the CRM user wants to ensure that all customer data is free of errors, and in case of PII the customers may have a legal right to demand that all information is correct.

Next, some mechanisms are important for secure elasticity:

- Secure data migration (E1) - confidential data must be securely moved between physical storage resources

Of the general security mechanisms shown in Figure 2, the CRM user will want to ensure that the provider employs:

- Access control mechanisms (AC1 and AC2) - to authenticate and authorize customers and to repel external attackers
- Logging (AU1) - CRM user wants to be able to determine after the fact if errors were made or breaches occurred.

The resulting security SLA with the required mechanisms is then shown in Figure 4.

6 DISCUSSION AND CONCLUSION

This paper presents a first step toward including security mechanisms in Cloud SLAs. The purpose is to increase trust in the Cloud by simplifying the process of designing security contracts for Cloud services, and in the next step, to enable services to be composed from a specific set of security requirements. The main advantage of our approach in its current form is its practical applicability, which we have tried to demonstrate with our simple case study.

One of the main unsolved challenges related to our approach is related to the difficulty of auditing that

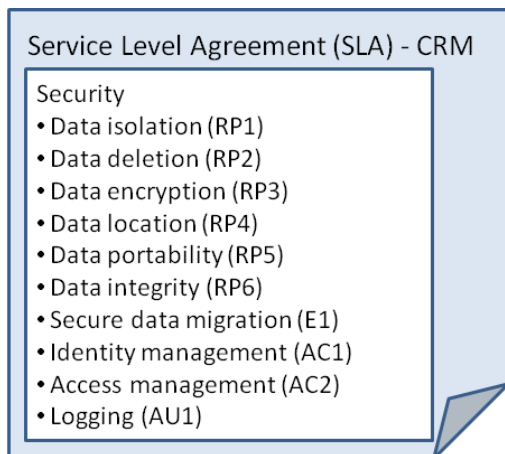


Figure 4: Security SLA for the CRM service.

security promises stated in the SLAs are being kept. Of the general security mechanisms shown in Figure 2, internal access control mechanisms, audit and incident management can typically be verified against a checklist according to ISO/IEC 27002 (27002, 2005). In the context of a SLA, the provider will provide yes/no answers to most of these items, whereas other items can be graded (how often is backup performed, is there a dedicated incident response team, etc.). Here it is possible to learn from the results produced by e.g. the CloudAudit⁴ association.

We are in the process of designing a management architecture based on our proposed framework. Its main tasks will be to compose Cloud services in accordance to the security requirements stated in a SLA. In addition, the management architecture will provide auditing mechanisms that will be used to verify that the contracted level of security is being obeyed. As a first step we will design and implement support for simple Cloud services, which require only a few security mechanisms, in order to evaluate the feasibility of our approach. The long-term vision is to establish a comprehensive architecture where SLAs including security, dependability and performance guarantees will form the basis for service composition.

7 ACKNOWLEDGEMENTS

This work has been supported by Telenor through the SINTEF-Telenor research agreement.

⁴<http://www.cloudaudit.org/>

REFERENCES

- 27002 (2005). Information technology – Security techniques – Code of practice for information security management.
- Casola, V., Mazzeo, A., Mazzocca, N., and Rak, M. (2006). A SLA evaluation methodology in Service Oriented Architectures. In Gollmann, D., Massacci, F., and Yautsiukhin, A., editors, *Quality of Protection*, volume 23 of *Advances in Information Security*, pages 119–130. Springer US.
- Christodorescu, M., Sailer, R., Schales, D. L., Sgandurra, D., and Zamboni, D. (2009). Cloud security is not (just) virtualization security: a short paper. In *Proceedings of the 2009 ACM workshop on Cloud computing security, CCSW '09*, pages 97–102, New York, NY, USA. ACM.
- de Chaves, S. A., Westphall, C. B., and Lamin, F. R. (2010). SLA Perspective in Security Management for Cloud Computing. In *Proceeding of the 2010 Sixth International Conference on Networking and Services*, pages 212–217. IEEE.
- European Network and Information Security Agency (ENISA) (2009). Cloud Computing: Benefits, risks and recommendations for information security.
- Frankova, G. and Yautsiukhin, A. (2007). Service and protection level agreements for business processes. In *Young Researchers Workshop on Service*.
- Heiser, J. and Nicolett, M. (2008). Assessing the Security Risks of Cloud Computing.
- Henning, R. R. (2000). Security service level agreements: quantifiable security for the enterprise? In *Proceedings of the 1999 workshop on New security paradigms, NSPW '99*, pages 54–60, New York, NY, USA. ACM.
- International Telecommunication Union (2008). Terms and Definitions Related to Quality of Service and Network Performance Including Dependability, ITUT E.800.
- Irvine, C. (2000). Quality of security service. In *Proc. ACM New Security Paradigms Workshop*, pages 91–99.
- Lindskog, S. (2005). *Modeling and tuning security from a quality of service perspective*. Chalmers University of Technology.
- Mell, P. and Grance, T. (2009). The NIST Definition of Cloud Computing, v.15.
- Payne, A. and Frow, P. (2005). A strategic framework for customer relationship management. *Journal of Marketing*, 69(4):167–176.
- Righi, R. R., Kreutz, D. L., and Westphall, C. B. (2006). Sec-mon: An architecture for monitoring and controlling security service level agreements. In *XI Workshop on Managing and Operating Networks and Services*.
- Verendel, V. (2009). Quantified security is a weak hypothesis: a critical survey of results and assumptions. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pages 37–50. ACM.