Disclaimer

This copy is a preprint of the article self-produced by the authors for personal archiviation. Use of this material is subject to the following copyright notice.

IEEE Copyright notice

Cloud Security: from Per-Provider to Per-Service Security SLAs

Alessandra De Benedictis*, Valentina Casola*, Massimiliano Rak† and Umberto Villano‡
* *Università di Napoli Federico II, DIETI, Napoli, Italy*
alessandra.debenedictis@unina.it, casolav@unina.it
† *Seconda Università di Napoli, DII, Aversa, Italy*
massimiliano.rak@unina2.it
‡ *Università del Sannio, DING, Benevento, Italy*
villano@unisannio.it

***Abstract*—Cloud Security is still considered one of the main factors inhibiting the diffusion of the Cloud Computing paradigm. Potential Cloud Service Customers (CSCs) do not trust delegating every kind of resources and data to external Cloud Service Providers (CSPs). The problem grows in complexity due to the increasing adoption of complex supply chains: CSPs that offer Sofware-as-a-Service (SaaS) cloud services often do not have their own data centers, but just acquire resources and services from other CSPs. This makes it hard, if not impossible, to ascribe the responsibility of a security incident. A possible solution is the adoption of Security Service Level Agreements (SLAs): CSPs should deliver services with an SLA that details each guarantee offered in terms of security, and CSCs should be able to compare offerings from different CSPs and verify that SLAs are respected during service life cycle. This paper shows how it is possible to build up a *per-service* Security SLA in a chain of cloud services, proposing a solution based on a security evaluation technique to compare different cloud service supply chains based on their Security SLAs.**

***Keywords*-Cloud, Cloud security, security SLA, comparing Cloud Service Providers**

I. Introduction

As outlined by recent ENISA documents [1], [2], [3], security is still considered an inhibiting factor for the diffusion of the cloud computing paradigm. As a matter of fact, prospective Cloud Service Customers (CSCs) typically do not trust external Cloud Service Providers (CSPs), and hesitate to delegate the management of their resources and data. Currently the problem is getting worse, due to the increasing adoption of complex supply chains, where services offered by a CSP are dependent on services (typically infrastructural ones) offered by other CSPs. For example, CSPs that offer Software-as-a-Service (SaaS) cloud services often rely on storage and compute resources hosted in data centers of external CSPs. In these situations, it is difficult to define who must grant the desired security level, or to ascertain the actual responsibilities in the case of a security incident.

Cloud Service Level Agreements (SLAs) [4] and, more specifically, security-oriented SLAs (Security SLAs) [5], [6], [7], can be a solution to the issues described above. Cloud Security SLAs have been adopted in research projects like SPECS [8], MUSA [9], SLAReady [10], ESCUDO [11], and supported by the EU Community [12]. Notwithstanding the wide research activity and the strong interest of the cloud customers, commercial cloud providers do not offer the desired Security SLAs. Currently cloud SLAs are essentially descriptions in natural language that focus only on few service terms (mostly on availability), completely ignoring all security-related aspects. Moreover, CSPs only offer SLAs with which they guarantee these service terms uniformly for all offered services to all customers, regardless of particular service characteristics or customers specific needs.

In this paper, we promote a *per-service* Security Service Level Agreement model, which enables both the provider and the customer to reach an agreement on the security features of each service instance being offered/leased. This model entails the use of a "tailored" SLA for each service, in that every customer can stipulate a (possibly) different SLA for each leased service. The feasibility of the proposed approach is demonstrated by its application in the context of the SPECS project [7], [13], whose aim is to develop a framework to automatically negotiate, monitor and enforce Security SLAs.

With the proposed approach, the end-user is provided with a well-defined set of security guarantees related to the acquired service, and, at the same time, all responsibilities related to possible security incidents occurring in the service supply chain are clearly assigned to the respective providers. Furthermore, it is also possible to compare different cloud service supply chains based on associated Security SLAs. The solution we propose is based on the REM technique (*Reference Evaluation Model* [14]), which has been already applied to provide a quantitative evaluation of the provided level of security [15].

The remainder of this paper is organized as follows: Section II deals with the state of the art on Security SLA and illustrates the Security SLA model proposed in the context of the SPECS project. Sections III and IV show how it is possible to write and evaluate a Security SLA by using per-provider information available to customers in open repositories. Section V describes how to obtain a per-service security SLA in the case of a chain of service invocations. The paper ends with the conclusions and with a discussion of future work (Section VI).

II. Cloud SLAs and security representation

Historically, the adoption of SLAs in the context of cloud computing was inspired by network systems and GRID systems, where they are widely used. However, it is a fact that exploiting SLAs in the cloud is more complex, due to the lack of a single reference technology (as Globus [16] in GRID) and of well-assessed standards. Moreover, the issue is complicated by the multiplicity of deployment and organizational models available for cloud computing.

The use of security-oriented SLAs to specify security requirements was first proposed by Henning [17] in 2000. Since then, it is still an open issue, due to the difficulties in representing security in a quantifiable way and in automating security best practices, commonly based on experts' experience. As discussed in [18], extending cloud SLAs to cover security aspects, allowing composition of cloud services from several service providers with a defined security level, is a challenging task. Extensive work has been recently done on Security SLAs with focus on clouds, in the context of academic, industry and government-driven initiatives. In 2011, ENISA published a report analyzing the use of security parameters in cloud SLAs [19]. This report was based on a survey of real-world CSP SLAs, and listed a set of common Security Level Objectives (SLOs). The report showed that, although security was considered by most respondents as a top concern and SLAs were often adopted by CSPs, they typically addressed only availability and other performance-related parameters, neglecting security-related ones. Following this initiative, the C-SIG SLA subgroup, an industry group supported by the European Commission, released in 2014 a set of SLA standardization guidelines [20] for CSPs. These guidelines provide definitions of the legal and technical terms used in SLAs, and identify SLOs specifically designed to achieve standardization for aspects of SLAs as secure data management and protection.

The identification of the security level objectives relevant for a service is related to the specification of the *security controls* that the provider is able to implement for that service. In order to promote the adoption of security best practices and aid the process of security management for enterprises, several standard initiatives have been launched in the last years, aiming at defining shared catalogs of security controls. Among these, the ENISA's Information Assurance Framework [21], released in 2009 and based on ISO 27001/2 standards and on industry best-practice requirements, was designed to help organizations assess the risk related to the adoption of cloud services and to compare different offers with respect to security properties. Other relevant initiatives are represented by the Cloud Control Matrix [22] released by the Cloud Security Alliance and the Control Framework proposed by NIST in its Special Publication sp800-53 [23].

Security SLAs and the automatic management of their life cycle were the focus of the *SPECS* EU FP7 project (ended in May 2016). In particular, SPECS was aimed at developing an open source framework of services and tools supporting the automatic management of the main phases of the Security SLA life cycle, namely *Negotiation*, *Implementation*, *Monitoring*, *Remediation* and *Renegotiation* [24]. The SPECS framework can be used to build secure cloud applications by enhancing existing (IaaS) cloud services with the deployment and configuration of ad-hoc security mechanisms and related monitoring systems, delivered as-a-service. The methodology adopted for building secure applications via the SPECS framework is fully described in [25], while [26] illustrates the SPECS approach to security monitoring based on SLAs. One of the main outcomes of the SPECS project is the introduced *Security SLA model*, which is the basis for the discussion carried out in this paper. The SPECS Security SLA model, described in detail in [27], is based on the WS-Agreement standard [28], which has been extended with provider-specific information and security-related concepts. In particular, in such model, security guarantees are specified in terms of the set of enforced standard security controls and of the (security) metrics associated to such controls that can be used to monitor their correct implementation.

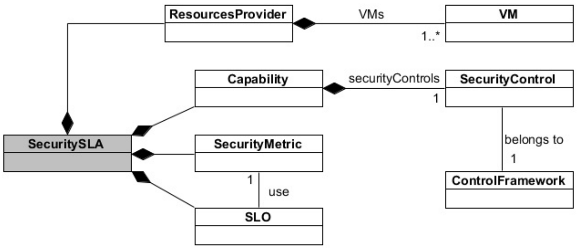

Figure 1. The SPECS Security SLA model: security-specific concepts

Figure 1 shows the concepts introduced to enable the per-service Security SLA approach and their mutual relationships. The *Resources Provider* concept models the origin of the infrastructure resources (i.e., the Virtual Machines - VMs) used to build the service covered by the SLA. This information is fundamental to keep track of the service supply chain and may be used, as discussed later (Section V), to evaluate its overall level of security.

For what regards security, security-related *declarations* are introduced in the form of *Security Capabilities*. Security capabilities are defined by NIST as "sets of mutually reinforcing security controls" [23], and are intended here as the security features that are offered on top of the service for which the agreement is built. Capabilities express security features according to common security best practices, and are meant to be understood even by customers with basic security skills. They are enforced by means of suitable software and/or hardware mechanisms, which are deployed by the provider either on the resources of the customer or on external resources. It is clear that, in order to enable the customer to verify that the capabilities declared by a provider are actually implemented (since they have signed a formal agreement), suitable monitoring functionalities must be provided.

For capabilities, what to measure is defined through the *Security Metrics* reported in the Security SLA. Security metrics represent measurable parameters associated with each declared capability and with the specific security mechanisms deployed for their enforcement. Security guarantees are expressed as constraints on the admissible values of declared security metrics (*SLOs*), and represent the security levels that the service customer requires and that the service provider accepts to offer.

In the next section, we will discuss the state of art for what regards the available security declarations of cloud providers, and then we will illustrate how the our model can be populated with the information currently available.

III. Security SLAs and Cloud Service Providers

The model proposed in Section II enables to manage detailed information about the guarantees offered by CSPs to their customers. However, as already mentioned, at the state of the art the information provided by CSPs on their security guarantees is very limited. So, it is reasonable to raise doubts about the actual applicability of the proposed SLA model.

To discuss this issue, we will consider one of the emerging approaches to assess security in the cloud, and demonstrate that it is possible to write and manage a Security SLA. As outlined in the introduction, even if Security SLAs are not explicitly adopted by CSPs, some information is now available from third party repositories, as the one proposed by the *CSA STAR program* [29], which provides security assessment through the definition of a set of security controls belonging to the CSA's CCM framework. Related to this, it is worth mentioning the *Consensus Assessments Initiative Questionnaire* (CAIQ) [30], which is the result of a providers' security self-assessment process, promoted by CSA and consisting in a set of *yes or no* assertions related to the implementation of specific security controls. It is worth noting that the security declarations obtained by the CSPs' self-assessment do not offer any concrete guarantee about their real enforcement. They do not represent a contract among customers and providers, do not constitute valid legal documents, but simply represent a public declaration. Moreover, they cannot be monitored, as they do not include any concrete security metric.

The above mentioned initiatives reflect the state of the art of cloud security and of the concrete security grants offered by CSPs. The positive aspect of the CAIQ and the STAR repository is that they represent a public repository of declarations that enables a CSC to perform a comparison among the security levels offered by each CSP. On the other hand, the negative aspect is that the CAIQ contains about 300 questions (categorized in controls and control domains) and therefore it is very difficult to analyse them for CSP comparison from the CSCs perspective.

Going back to the model presented in the previous section, the CAIQ controls can be used to build the declarative part of a Security SLA, represented by the set of provided security capabilities. Our approach consists in including in the Security SLA of a CSP every control for which it has provided a *yes* reply in the CAIQ. Note that we consider a security control as actually implemented by a CSP if and only if it replied *yes* to all the related questions in the CAIQ. Unfortunately, this process cannot be completely automatized, since many CSPs did not provide a *yes* or *not* reply to questions, using instead a free-text description. When this happens, a security expert has to evaluate the CAIQ replies and *tune the response* according to his own policies.

The above process allows to build a Security SLA for each provider that has been interviewed with the CAIQ. This SLA represents the security provision/offer of the CSP. In the next section, we will show how it is possible for a CSC to compare different providers based on such SLAs, according to several criteria. In Section V, we will discuss how to compose Security SLAs from different providers involved in the supply chain of a specific service (*per-service* SLA), and how to include in this composition possible additional security guarantees that CSPs may be willing to offer.

IV. Comparing per-Provider CAIQ-based Security SLAs

As illustrated in the previous section, according to our approach Security SLAs of CSPs are built starting from the public information available in the CAIQ. In the following, we illustrate how on the basis of this representation a CSC can compare different offerings (i.e., different Security SLAs). The methodology adopted to compare Security SLAs is the *Reference Evaluation Model* (REM) [14], which consists of the following elements:

- *Formalization* is the formal representation (in XML) of the CSPs' security offers. In our case, this corresponds to the Security SLA XML representation built from the CAIQ (as discussed in the previous section).
- *Technique* represents the evaluation technique that is applied to compare offers. It is based on the euclidean distance among different CAIQ answers.
- *Reference Levels* are the sets of CAIQ results chosen as references (i.e., a CAIQ with all replies set to *yes*, or a CAIQ with all replies set to *no*), which represent different security levels.

Providing a detailed description of the methodology is out of the scope of this paper. The interested reader is referred to [14] for further details. However, to make the discussion self-consistent and easy to follow, we will recall some elements of methodology when needed.

REM assumes that Security SLAs are represented by means of a tree structure. Hence, we modeled the CAIQ (that is the basis for our SLA representation) as shown in Figure 2. The tree in the figure is rooted on the CAIQ of the CSP; the first children level represents the Control Groups, the second children level reports the security Controls, and the leafs are associated to the questions of the CAIQ.

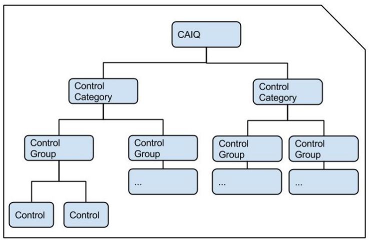

Figure 2. The SLA hierarchy representing a CAIQ as a tree

The REM methodology consists of three phases, namely (i) *Policy Structuring*, (ii) *Policy Formalization* and (iii) *Policy Evaluation*.

The goal of the *Structuring* phase is to associate an enumerative and ordered data type K_i to the n leaves-provisions of the tree. In our case, CAIQ assumes that all the leaves may assume only three values: *yes*, *no* and *N/A* (Not Available). Hence, the enumerative type is simply defined as K=$\{yes,\ no,\ N/A\}$. Values are ordered from *yes* to *N/A*, with $yes > no$ and $no > N/A$. This is based on the reasonable assumption that a completely missing answer to a CAIQ question (*N/A* value) means that the security controls has not been addressed at all, while a *no* answer represents a decision after an evaluation of the issue. However, these decisions can be adapted during the evaluation steps, if needed.

According to the above model, the CAIQ space P is defined as $P = K^n$, which means that it is a vector containing the n replies.

In the *Formalization* phase, the CAIQ space P is turned into a homogeneous space PS. This transformation is accomplished by a normalization and clustering process, which allows to associate a *Local Security Level* (LSL) to each security offer. After that, the provisions may be compared by comparing their LSLs. In the CAIQ, the clustering process is simple, as all replies may assume only three possible values (yes-3, no-1, n/a-0).

The goal of the *Evaluation* phase is to pre-process the PS vector of LSLs, and evaluate the so-called *Global Security Level* L_{Px} associated to the provision P_x. The GSL has been defined on the basis of an Euclidean distance among matrices and some reference levels:

$$L_{Px} = \begin{cases} L_0 & iff d_{x0} \leq d_{10} \\ L_1 & iff d_{10} < d_{x0} < d_{20} \\ L_2 & iff d_{20} < d_{x0} < d_{30} \\ L_3 & iff d_{30} < d_{x0} < d_{40} \\ L_4 & iff d_{40} \leq d_{x0} \end{cases}$$

where $d_{i,0}$ are the distances among the references and the origin of the metric space (denoted as $\emptyset$), as illustrated in [14]. This function gives a numerical result to the security; the GSL is a measure of the security provided by the CSP infrastructure according to its CAIQ, and so it can be easily used to compare different providers.

We have developed a demonstrator application, available at http://apps.specs-project.eu:8080/specs-app-SecurityReasoner, through which it is possible to upload different CAIQs and compare them against specific requirements, with the possibility of assigning different weights to the available control groups.

V. From per-Provider to per-Service Security SLAs

In Section III we have shown how it is possible to build Security SLAs for existing services by exploiting an open repository of information, and in Section IV we have discussed how to compare CSPs based on their security declarations. In this section, we want to investigate what happens to the Security SLAs of services in the presence of a service supply chain composed of more than one CSP, also considering the case when some of the CSPs of the chain offer specific security guarantees. Stated another way, we will show how is it possible to compose the offered Security SLAs and to compare different available offers.

An example may help to identify the problem to be tackled. Let us consider a CSP (CSP1) offering to its customers a *reliable* web server solution (in a Platform-as-a-Service fashion). The reliable web server offered by CSP1 is built on top of an Infrastructure-as-a-Service service (a set of VMs) offered by CSP2.

Let us assume that CSP1 offers security mechanisms integrated in its PaaS solution, which increase the security level of the service offered to the CSC. Moreover let us assume that, in order to attract more customers, CSP1 is willing to offer a concrete Security SLA for its reliable web server solution. This includes monitorable security metrics and SLOs. While the security mechanisms implemented by CSP1 are under its control, no guarantees can be provided by CSP1 for the infrastructure services acquired from CSP2. In practice, CSP1 may result less secure than believed.

The SPECS Security SLA model we propose makes it possible to address this issue thanks to the declaration of security capabilities and to the introduction of resource providers information. In our example, the CSP1 Security SLA will include the declaration of a *Web Resiliency* capability (built by CSP1 by means of *ad-hoc* mechanisms) and the information on CSP2 and on the number and type of VMs acquired from it. Hence, CSP1 offers to its own customers a *per-service* Security SLA, which clearly states the security guarantees for which CSP1 is responsible and the metrics to be used to monitor the characteristics of the delivered service.

Based on this Security SLA, it is possible to enrich the information already present in the CAIQ repository about CSP1 and CSP2 with the additional information related to the new security features provided by CSP2, information that is present in its Security SLA. It should be noted that the reliable web server offered by CSP1 can thus be compared with other services, exploiting the methodology presented in Section IV.

In the first step of the described process, the Security SLA is processed to obtain an equivalent description (an SLA tree hierarchy) to be used for evaluation. Then, from the SLA tree hierarchy it is possible to identify and retrieve the CAIQ replies associated with the involved CSPs with respect to controls present in the hierarchy. In the subsequent step, the declared capabilities are extracted from the Security SLA, in order to know which are the additional controls that it is possible to enforce through the implemented security mechanisms, adding them to the SLA tree hierarchy. The final result, shown in Figure 3, is that it is possible to generate an *enhanced* CAIQ, which does not simply contain the basic replies of CPSs, but also possible additional security guarantees offered by specific providers.

Finally, the REM evaluation technique, illustrated in the previous section, can be used to evaluate and compare the enhanced CAIQs associated to the different Security SLAs. It is worth pointing out that the customer will be aware that the offers of CSP1 are enriched with a clear security responsibility (but only for what it offers on-premise), while CSP2 provides only public declarations.

VI. Conclusions

In this paper, we have tackled the problem of evaluating and comparing the level of security offered by cloud providers in the presence of complex service supply chains, involving the acquisition of resources from different CSPs. Our approach founds on the adoption of a novel model of Security Service Level Agreement (Security SLA), developed in the context of the European project SPECS, and meant to define the security guarantees offered on each specific service (*per-service* SLA).

We have illustrated the process adopted to build *per-service* Security SLAs starting from the security declarations of existing providers that are currently available in public repositories (STAR repository). According to the proposed process, the information on the security controls put in place by providers (derived from the answers given by providers to the CSA's CAIQ) is used to fill the declarative section of their respective Security SLAs (*per-provider* SLAs). These sections identify the security features provided by the CSPs, without specifying any means to verify their actual delivery.

Subsequently, in the presence of a complex supply chain in which existing cloud services (whose security features are specified through the CAIQ results) are combined with additional security features offered *as-a-service* by one or more providers, a Security SLA for the service is obtained. This combined Security SLA is built by generating an *enhanced* CAIQ result, including all the replies given by the CSPs involved in the chain, plus *ad-hoc* updated entries depending on the additional security features offered on top of the target service. Moreover, this SLA can possibly contain security metrics useful for monitoring the additional security capabilities introduced in the supply chain. Finally, we have shown how to compare the Security SLAs associated with different service supply chains by adopting the REM methodology.

The combined *per-service* Security SLA is currently obtained by simply putting together all the controls declared by CSPs in the supply chain. In future work, we plan to investigate the relationship among security controls in order to be able to identify which controls are actually guaranteed, on top of the target service as a whole.

Acknowledgment

This research is partially supported by the grant FP7-ICT-2013-11-610795 (SPECS) and H2020-ICT-07-2014-644429 (MUSA).

References

[1] ENISA, "Procure Secure. A guide to monitoring of security service levels in cloud contracts," April 2012.

[2] M. Dekker, "Critical cloud computing a ciip perspective on cloud computing services," , ENISA, Tech. Rep., 2012.

[3] D. Catteddu, "Security & resilience in governmental clouds," , CSA, Tech. Rep., 2011.

[4] CSCC, "The cscc practical guide to cloud service level agreements," http://www.cloudstandardscustomercouncil.org/webSLA-download.htm, CSCC, Tech. Rep., 2012.

[5] G. H. Marnix Dekker, "Survey and analysis of security parameters in cloud slas across the european public sector," 2011. [Online]. Available: http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/survey-and-analysis-of-security-parameters-in-cloud-slas-across-the-europe

[6] A. Pannetrat, G. Hogben, S. Katopodis, G. Spanoudakis, and C. Cazorla, "D2.1: Security-aware sla specification language and cloud security dependency model. technical report, certification infrastructure for multi-layer cloud services (cumulus)." 2013.

[7] M. Rak, N. Suri, J. Luna, D. Petcu, V. Casola, and U. Villano, "Security as a service using an sla-based approach via specs," in *Proceedings of IEEE CloudCom Conference 2013*, IEEE, Ed., 2013.

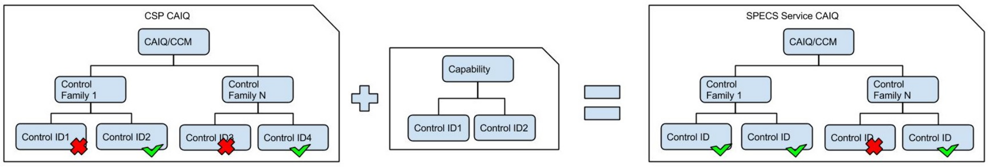

Figure 3. Generating the enhanced CAIQ

[8] SPECS Consortium, "The specs project web site," http://specs-project.eu/, 2013.

[9] MUSA Consortium, "The musa project web site," http://musa-project.eu/, 2015.

[10] SLA Ready Consortium, "The sla ready project web site," http://www.sla-ready.eu/, 2015.

[11] ESCUDO-CLOUD Consortium, "The escudo-cloud project web site," http://www.escudocloud.eu/, 2015.

[12] Cloud Select Industry Group, "Cloud Service Level Agreement Standardisation Guidelines," Commissions European Cloud Strategy, Tech. Rep., 2014. [Online]. Available: https://ec.europa.eu/digital-agenda/en/news/cloud-service-level-agreement-standardisation-guidelines

[13] V. Casola, A. De Benedictis, M. Rak, and U. Villano, "Preliminary Design of a Platform-as-a-Service to Provide Security in Cloud," in *CLOSER 2014 - Proceedings of the 4th International Conference on Cloud Computing and Services Science, Barcelona, Spain, April 3-5, 2014.*, 2014, pp. 752–757.

[14] V. Casola, A. Mazzeo, N. Mazzocca, and V. Vittorini, "A Policy-Based Methodology for Security Evaluation: A Security Metric for Public Key Infrastructures," *Journal of Computer Security*, vol. 15, no. 2, pp. 197–229, 2007.

[15] V. Casola, M. Rak, and G. Alfieri, "A cloud application for security service level agreement evaluation," in *Proceedings of the 4th International Conference on Cloud Computing and Services Science*, 2014, pp. 299–307.

[16] I. Foster, *Network and Parallel Computing: IFIP International Conference, NPC 2005, Beijing, China, November 30 - December 3, 2005. Proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, ch. Globus Toolkit Version 4: Software for Service-Oriented Systems, pp. 2–13. [Online]. Available: http://dx.doi.org/10.1007/11577188_2

[17] R. R. Henning, "Security service level agreements: Quantifiable security for the enterprise?" in *Proceedings of the 1999 Workshop on New Security Paradigms*, ser. NSPW '99. New York, NY, USA: ACM, 2000, pp. 54–60. [Online]. Available: http://doi.acm.org/10.1145/335169.335194

[18] K. Bernsmed, M. Jaatun, P. Meland, and A. Undheim, "Security slas for federated cloud services," in *Availability, Reliability and Security (ARES), 2011 Sixth International Conference on*, Aug 2011, pp. 202–209.

[19] M. Dekker and G. Hogben, "Survey and Analysis of Security Parameters in Cloud SLAs Across the European Public Sector," ENISA, Tech. Rep., December 2011.

[20] European Commission – C-SIG (Cloud Select Industry Group) subgroup, "Cloud Service Level Agreement Standardisation Guidelines," Tech. Rep., June 26 2014.

[21] "Information Assurance Framework," ENISA, Tech. Rep., November 2009.

[22] Cloud Security Alliance, "Cloud Control Matrix v3.0," https://cloudsecurityalliance.org/download/cloud-controls-matrix-v3/.

[23] NIST, "SP 800-53 Rev 4: Recommended Security and Privacy Controls for Federal Information Systems and Organizations," National Institute of Standards and Technology, Tech. Rep., 2013. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf

[24] A. De Benedictis, M. Rak, M. Turtur, and U. Villano, "Rest-based sla management for cloud applications," in *Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), 2015 IEEE 24th International Conference on*, June 2015, pp. 93–98.

[25] M. Rak, U. Villano, V. Casola, and A. De Benedictis, "Sla-based secure cloud application development: the specs framework," in *Symbolic and Numeric Algorithms for Scientific Computing, 2015 17th International Symposium on*, 2015, p. To appear.

[26] V. Casola, A. De Benedictis, and M. Rak, "Security monitoring in the cloud: An sla-based approach," in *10th International Conference on Availability, Reliability and Security, ARES 2015, Toulouse, France, August 24-27, 2015*, 2015, pp. 749–755.

[27] V. Casola, A. De Benedictis, M.Rak, J. Modic, and M. Erascu, "Automatically enforcing security slas in the cloud," *IEEE Transactions on Services Computing (PrePrints)*, 2016.

[28] A. Andreieux, "Web services agreement specification," 2007. [Online]. Available: https://www.ogf.org/documents/GFD.107.pdf

[29] Cloud Security Alliance, "Security, Trust & Assurance Registry (STAR)," https://cloudsecurityalliance.org/star/, 2011.

[30] ——, "Consensus Assessment Initiative Questionnaire," https://cloudsecurityalliance.org/group/consensus-assessments/, 2011.