

Quantitative Assessment of Cloud Security Level Agreements A Case Study

Jesus Luna, Hamza Ghani, Tsvetoslava Vateva and Neeraj Suri

Department of Computer Science, Technische Universität Darmstadt, Darmstadt, Germany

Keywords: Cloud Security, Security Assessment, Security Benchmarks, Security Level Agreements, Security Metrics.

Abstract: The users of Cloud Service Providers (CSP) often motivate their choice of providers based on criteria such as the offered service level agreements (SLA) and costs, and also recently based on security aspects (i.e., due to regulatory compliance). Unfortunately, it is quite uncommon for a CSP to specify the security levels associated with their services, hence impeding users from making security relevant informed decisions. Consequently, while the many economic and technological advantages of Cloud computing are apparent, the migration of key sector applications has been limited, in part, due to the lack of security assurance on the CSP. In order to achieve this assurance and create trustworthy Cloud ecosystems, it is desirable to develop metrics and techniques to compare, aggregate, negotiate and predict the trade-offs (features, problems and the economics) of security. This paper contributes with a quantitative security assessment case study using the CSP information found on the Cloud Security Alliance's Security, Trust & Assurance Registry (CSA STAR). Our security assessment rests on the notion of Cloud Security Level Agreements — SecLA — and, a novel set of security metrics used to quantitatively compare SecLAs.

1 INTRODUCTION

Despite the pervasive nature of Cloud technologies and their advocated economic/technological advantages, the migration of applications has been limited, in part, due to the lack of *security assurance* by the CSP. This lack of assurance, along with the current paucity of techniques to quantify security, often results in users being unable to assess the security of the CSP they are paying for. Despite the assumption that a given Cloud provider “seems” secure, is it actually “secure enough” for my needs? Is my personal data more secure today than before? How do I compare against other providers with regards to security? These questions have been raised in the security metrics area, including institutions such as ENISA¹ (Trimintzios, P., 2011), NIST² (Jansen, W., 2010) and CIS³ (Boyle K., et.al., 2010). The stated belief is that well-designed security metrics and quantitative security assessment techniques will both support the achievement of assurance in CSPs and also motivate the creation of trustworthy Cloud ecosystems.

Fortunately, security assurance in Cloud comput-

¹European Network and Information Security Agency.

²National Institute of Standards and Technology.

³Center for Internet Security.




Figure 1: Representation of the quantitative Cloud security assessment presented in this paper. Dashed shapes indicate the paper contributions.

ing has recently initiated promising steps. In particular, the Cloud community has identified that specifying security in Service Level Agreements (termed as “Security Level Agreements” or SecLA over this paper) to be useful to model and assess the security being offered by a CSP (cf. (Dekker, M. and Hogben, G., 2011) and (Cloud Security Alliance, 2011c)).

Despite the potential advantages related with the design and use of Cloud SecLAs, the security metrics and techniques to quantitatively reason about Cloud SecLAs are still lacking. To address this gap, this

paper develops a quantitative security assessment approach considering a case study using the CSP information found on the STAR repository (Cloud Security Alliance, 2011c).

Our proposed quantitative security assessment approach is illustrated in Figure 1 and, is based on the notion of Cloud SecLAs, with the contributions being:

1. A novel set of security metrics used to compare, *benchmark* and assess the security compliance of the evaluated Cloud SecLAs.
2. A real-world case study of the proposed security assessment using the CSP information found on the STAR repository.

Our results aim at impacting the Cloud ecosystem in terms of security and transparency, thus facilitating service compliance and removing uncertainty from CSPs interactions (ultimately helping them to decrease transaction costs).

The paper is organized as follows: Section 2 motivates the notion of Cloud Security Level Agreements; Section 3 presents the proposed quantitative security assessment techniques, and Section 4 discusses the results of the contributed security assessment using the CSP information found on the STAR repository. Section 5 reviews relevant existing approaches related with our research.

2 CLOUD SECURITY LEVEL AGREEMENTS

The notion of SecLAs currently exists in varied dimensions (please refer to Section 5 for related state of the art) and the Cloud is not an exception. The use of Cloud SecLAs has the potential to provide tangible benefits to CSPs especially associated with improved security administration and management practices, thus allowing for more transparency to end users. It is clear that the definition of SecLAs forces a stakeholder to think about security. However end users can also benefit from SecLAs by understanding the costs and benefits associated with this new service model. Nowadays it is being realized that on the one hand SecLAs will provide service-based assurance, but on the other hand it is clear that SecLAs are not intended to replace electronic assurance mechanisms for security policy enforcement (Henning, R., 1999).

The importance of Cloud SecLAs has also been recognized by ENISA: the development of template contracts and service level agreements is being highlighted as one of the areas to be addressed in the European Cloud computing strategy. In a recent survey

(Dekker, M. and Hogben, G., 2011) ENISA highlights that many Cloud customers often do not monitor security aspects of their contracted SLA on a continuous basis. This implies that customers are left unaware about many important security aspects related to their services. The risk is that they find out about failing security measures only following a security breach. The survey data shows that while SLAs are often used, and availability is often addressed in these SLAs, security parameters are less well covered.

As mentioned in our previous work (Luna, J., et.al., 2011) and shown in Figure 2, the notion of Cloud SecLAs is essential in order to quantitatively assess the security level offered by a CSP. Cloud SecLAs usually model the CSP security at the *service level* and, are based on either a set of security requirements (e.g., for compliance reasons) or some kind of preliminary security threat analysis. The result will be a collection of security statements (also called “security provisions”) related with the CSP’s technical and organizational security, therefore ranging from provisions like “Backup periodicity” to “Independent audits frequency”. This final set of security provisions will become the Cloud SecLA, just as shown in Figure 2 and discussed in works like (Bernsmed, K., et.al., 2011) and (Almorsy, M., et.al., 2011).

Apart from the challenges related with the creation of SecLAs in real Cloud deployments, the current paucity of techniques to quantitatively assessing them has proven to be part of the obstacles in using SecLAs. To achieve the security assessment part shown in Figure 2, this paper presents techniques to quantitatively reason about Cloud SecLAs.

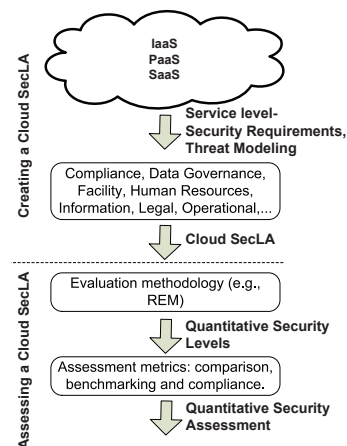


Figure 2: Workflow to quantitatively assess Security Level Agreements in Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS) CSPs.

3 SECURITY ASSESSMENT

Based on the notion of Cloud SecLA presented in Section 2, it is possible to derive a set of security metrics to quantitatively assess the CSP from three different perspectives: Comparison (Section 3.2), Benchmarking (Section 3.3) and Compliance (Section 3.4).

However, before actually introducing the contributed assessment metrics it is helpful to have a mechanism to quantitatively reason about SecLAs. Our approach utilizes the Reference Evaluation Methodology — REM — originally proposed in (Casola V., et.al., 2005) as a technique to quantitatively evaluate security policies. The basics related with the use of the REM technique are presented next.

3.1 The Reference Evaluation Methodology at Glimpse

In its basic form, REM (*a*) considers a set of security provisions⁴ to evaluate, (*b*) formalizes it to facilitate the further evaluation step over an homogeneous metric space, (*c*) uses a set of reference levels (known as *Local Security Levels* or *LSL*) to apply a distance criterion, and (*d*) finally obtains a number (also called *Global Security Level* or *GSL*) that corresponds to the policy's security level.

For the purposes of this paper, only a couple of REM-related concepts are presented and the interested readers are referred to (Casola V., et.al., 2005) for further details.

The first concept is the formal representation of any security policy by a $n \times m$ matrix, whose n rows represent single provisions with a maximum of m possible LSLs. For example, if the LSL associated to a Cloud Storage Provider's provision called "File System Encryption" is 3, then the corresponding vector⁵ will be $(1, 1, 1, 0)$. The REM concept of LSL is congruent with the notion of security ranges, as presented in (Irvine, C. and Levin, T., 2001).

The second REM-concept is a criteria used to quantify GSLs, and defined as the Euclidean distance⁶ among whatever pair of REM-matrices (A, B) . Just as in Equation 1, the GSL is defined as the square root of the matrix trace of $((A - B)(A - B)^T)$, where $(A - B)^T$ is the conjugate transpose.

$$GSL(A, B) = d(A, B) = \sqrt{Tr((A - B)(A - B)^T)} \quad (1)$$

⁴In REM terminology, a security provision is a security statement in the form {attribute, value}.

⁵The LSL is usually known as the *L1-Norm* (Weisstein, W., 2011b) of this vector.

⁶Also known as the Frobenius-Norm (Weisstein, W., 2011a).

We consider utilizing REM primarily for the flexibility it offers to model and quantitatively evaluate most classes of security policies (e.g., Certificate Policies (Casola, V. et.al., 2007), (Casola, V., et.al., 2007) and security policies associated with Desktop Grids (Luna, J. et.al., 2008)). However, as discussed later in Section 6, our security assessment can be easily extended to use other SecLA quantification techniques.

As a stand-alone security evaluation methodology, REM does not provide any additional metric to quantitatively reason about the LSL and GSL associated with each Cloud SecLA. To bridge this gap, our research develops a novel set of metrics to perform the security assessment of a CSP based on the REM-quantification of its SecLA, as depicted in Figure 2 and in Section 3.2.

3.2 Comparing Security Levels

The metrics introduced next are used to quantitatively compare the GSL (as computed with Equation 1) of two or more Cloud SecLAs, by determining where in the "SecLA metric space — MS —" are. The MS used in this section and graphically shown in Figure 3, is delimited by the following two SecLA values:

- $SecLA(\phi) = GSL(\phi, \phi)$, where ϕ is the origin of the metric space or in other words, a Cloud SecLA where all LSLs equal zero (i.e., are represented by the vector $(0, 0, 0, 0)$).
- $SecLA(max) = GSL(max, \phi)$ is the maximum allowable GSL for any Cloud SecLA under evaluation, and max is the SecLA where all LSLs are represented by $(1, 1, 1, 1)$.

Notice that the SecLA metric space mentioned above is able to contain whatever Cloud SecLA that uses a maximum of four LSLs.

Definition 1. *The following metric computes the rate among the overall security level of the CSP under analysis and, the maximum achievable security level (i.e., the total length of the MS shown in Figure 3):*

$$RateMS(SecLA_{CSP}) = \frac{GSL(SecLA_{CSP}, \phi)}{GSL(max, \phi)}$$

Where:

- $SecLA_{CSP}$ is the SecLA of the CSP to assess.
- $GSL(CSP, \phi)$ is the GSL of the CSP as computed with Equation 1 and, taking as reference level the origin of MS.
- $GSL(max, \phi)$ is the length of the defined MS.

From Definition 1, it can be observed that the closer $RateMS$ to 1 is, the more secure the CSP is.




Figure 3: Metric Space (MS) for quantitatively comparing Cloud SecLAs.

Because in practice the overall quantification of the Cloud SecLA may “hide” some details related with the security level of individual provisions, we propose the metric shown in Definition 2 to perform a more fine-grained comparison.

Definition 2. To quantitatively compare a CSP_1 with respect to second CSP_2 , we proceed to split their SecLAs into a couple of subsets, just as defined next:

$$\Delta_+(SecLA_{CSP1}, SecLA_{CSP2}) = \frac{GSL(SecLA_{+CSP1}, SecLA_{+CSP2})}{GSL(SecLA_{+CSP2}, \phi)}$$

$$\Delta_-(SecLA_{CSP1}, SecLA_{CSP2}) = \frac{GSL(SecLA_{-CSP1}, SecLA_{-CSP2})}{GSL(SecLA_{-CSP2}, \phi)}$$

Where:

- $\Delta_+(SecLA_{CSP1}, SecLA_{CSP2})$ is the subset of security provisions from CSP_1 with a security level greater or equal than CSP_2 .
- Analogously, $\Delta_-(SecLA_{CSP1}, SecLA_{CSP2})$ is the subset of security provisions from CSP_1 with a security level smaller than CSP_2 .

From Definition 2, the closer both Δ_+ and Δ_- are to 0, the more similar CSP_1 to CSP_2 will be.

3.3 Benchmarking Cloud SecLAs

The notion of *benchmarking* used in Definition 3 is to count *how many individual provisions* of Cloud SecLA are above/below of a predefined baseline (e.g., user-defined or standard-based). This metric is inspired by CIS’ security benchmark tool (Center for Internet Security, 2009) and, uses the concept of LSL explained in Section 3.1.

Definition 3. If two Cloud SecLAs have the same cardinality, then the “Benchmarking Score” of a CSP’s $SecLA_{CSP}$ relative to a reference $SecLA_{REF}$ is defined as:

$$BenchScore(SecLA_{CSP}, SecLA_{REF}) = \frac{|\bigcup LSL_{CSP}|}{|SecLA_{CSP}|}$$

Where:

- $|\bigcup LSL_{CSP,n}|$ is the number of security provisions from $SecLA_{CSP}$, such that $LSL_{CSP} \geq LSL_{REF}$.
- $|SecLA_{CSP}|$ represents the SecLA’s cardinality (i.e., the total number of security provisions from either $SecLA_{CSP}$ or $SecLA_{REF}$).

From Definition 3, the closer *BenchScore* to 1 is, the best $SecLA_{CSP}$ has the baseline $SecLA_{REF}$ fulfilled. However, if $BenchScore < 1$ then the CSP is only *partially compliant*, whereas if $BenchScore = 0$ then the CSP is *not compliant* at all with $SecLA_{REF}$.

3.4 Compliance Assessment

This metric is a case of the Benchmarking Score metric presented in the previous section. The Compliance Index metric shown next was designed to assess if a CSP’s SecLA fulfills (or not) some specific security compliance criteria.

Definition 4. Using the same notation from Definition 3, the quantitative “Compliance Index” (*CompIndex*) of a CSP relative to a reference criteria is defined as:

$$CompIndex(SecLA_{CSP}, SecLA_{REF}) = \begin{cases} 1 & \text{if } LSL_{CSP,n} \geq LSL_{REF,n} \\ & \forall LSL_{CSP,n} \in SecLA_{CSP} \wedge \\ & \forall LSL_{REF,n} \in SecLA_{REF} \\ \frac{|\bigcup LSL_{CSP}|}{|SecLA_{CSP}|} & \text{otherwise} \end{cases}$$

From Definition 4, if $CompIndex(SecLA_{CSP}, SecLA_{REF}) = 1$ then it means that the CSP’s SecLA is *fully compliant* with the reference criteria $SecLA_{REF}$.

4 EVALUATION: A PRACTICAL SECURITY ASSESSMENT STUDY

A well-known challenge in the security metrics field is related with the empirical validation of metrics and frameworks (Verendel, V., 2009). Fortunately, the Cloud Security Alliance (CSA) has recently released the initial version of their “Security, Trust & Assurance Registry” (STAR (Cloud Security Alliance, 2011c)), a publicly available repository that documents the security controls provided by CSPs worldwide. Currently, STAR contains the “Consensus Assessments Initiative Questionnaire” (CAIQ (Cloud Security Alliance, 2011b)) reports, which provides industry-accepted ways to document what security controls exist in Cloud offerings. The CAIQ contains a set of over 160 questions a Cloud consumer and Cloud auditor may wish to ask a CSP.

The remainder of this section presents the results obtained from applying the metrics (of Section 3), to the CAIQ reports stored in the STAR repository. Please notice that (i) at the time of writing this paper STAR contained only the reports of three CSPs and, (ii) due to STAR’s usage restrictions we have anonymized the identity of the CSPs under analysis.

4.1 Testbed

To perform in a semi-automated manner the contributed CSP’s security assessment, we have implemented an initial prototype of the architecture shown in Figure 4.

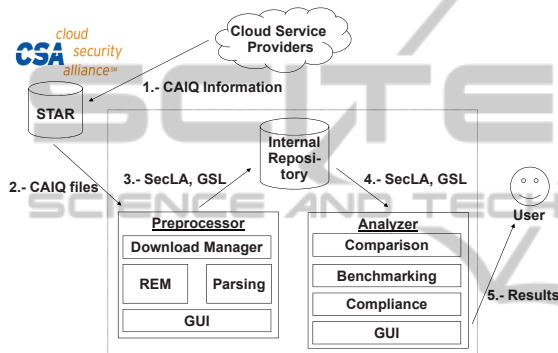


Figure 4: Testbed to perform the CSP’s security assessment.

At the core of our testbed are two building blocks with the following functionality:

1. **Preprocessor:** after a CSP has uploaded its CAIQ report to STAR (Step 1 in Figure 4) the *Download Manager* will retrieve it (Step 2). Once downloaded, the CAIQ report is parsed into the Internal Repository shown in Figure 4 via the *Parsing* module. Using a local *REM* implementation, the parsed Cloud SecLA is quantitatively evaluated and the resulting *GSL/LSL* stored into an *Internal Repository* (Step 3). In the current version of our prototype, both the downloading and parsing processes are manually performed via the *GUI* component.
2. **Analyzer:** this module will download from the Internal Repository a set of Cloud SecLAs to assess (Step 4). Depending on the metrics selected by the User (via the Analyzer’s *GUI*) this module will invoke one or more of the *Comparison*, *Benchmarking* and *Compliance* components. The results of the security assessment will be graphically displayed via the *GUI* (Step 5).

Using on a working prototype of the testbed explained above, we show in Section 4.3 the results of

the CSPs’ quantitative security assessment using the STAR reports.

4.2 Setting Up the SecLAs

Based on the CAIQ responses of the CSP under analysis, we created for the purposes of our security assessment the following two sets of Cloud SecLAs for each one of them:

1. **CAIQ:** this is an initial set of three SecLAs containing all the 171 security provisions from the different CAIQ’s “control groups” and common to all the CSPs under evaluation. These 171 provisions are distributed in the following way: Compliance (CO) - 14, Data Governance (DG) - 15, Facility Security (FS) - 9, Human Resources Security (HR) - 4, Information Security (IS) - 71, Legal (LG) - 2, Operations Management (OP) - 5, Risk Management (RI) - 12, Release Management (RM) - 5, Resilience (RS) - 11 and Security Architecture (SA) - 23). These provisions had a qualitative “YES/NO” answer in the CAIQ, but there was an additional “Comments” field that allowed the CSP to provide further details.
2. **CAIQ+:** after analyzing the CSP’s answers given on the CAIQ’s “Comments” field, we designed this additional set with 29 security metrics that were used to create another three SecLAs (in this case with $171 + 29 = 200$ security provisions). Table 1 shows some of these newly proposed metrics.

Finally, in order to establish our Metric Space (cf. Section 3 and Figure 3) we created an additional SecLA, where all security provisions were set to their maximum value ($LSL = 4$).

4.3 Results

Our first test consisted in applying the “RateMS” metrics (introduced in Section 3.2) to the two sets of SecLAs described in Section 4.2. Figure 5 show our obtained results for the three CSPs in STAR. It is interesting to notice that despite CSP_1 had the highest security level according to the CAIQ SecLA, once we have considered the additional set of metrics (the CAIQ+ SecLA) then it was CSP_3 the one with the better security level associated to its Cloud SecLA. The reason for this can be easily obtained after reading both CAIQ reports: on the one hand, CSP_1 replied “YES” to almost all the questions on the CAIQ but their “Comments” were so generic that it was not possible to obtain further quantitative data to evaluate the additional set of 29 metrics on a reliable way.

Table 1: Some of the proposed security metrics for the CAIQ, based on the CSP’s responses found on the STAR repository

CID	Consensus Assessment Questions	Possible values
CO-01.Q1	Supported Audit Assertions formats	None, Proprietary, Only one, More than one
CO-02.Q1	Access requirements to third party audit reports	Not specified, Only to customers, NDA Required, NDA and written justification required
CO-02.Q2	Frequency of the network penetration tests	Not specified, Once per-Year, Twice per-Year, Monthly
IS-19.Q2	Tenants’ encryption keys management	Not specified, Not supported, Provider only, Shared management, User only

On the other hand, CSP_3 ’s answers to the original CAIQ contained more “NO” values than CSP_1 ’s but its “Comments” were by far much richer in information to quantify the corresponding CAIQ+ SecLA. This demonstrates that disclosing a greater level of detail on the SecLA is beneficial for the CSP. Also notice that in all the cases the security levels associated with the CAIQ+ SecLAs were lower than those obtained for the CAIQ SecLAs, this was in part because of the difference between both metric spaces ($MS_{CAIQ} = 26.153$ whereas $MS_{CAIQ+} = 28.284$).

curity level (Δ_- metric) was by far larger also for both CSP_1 and CSP_2 (approximately 86% for the CAIQ subset and, between 75% — 79% for the CAIQ+ subset). Once again, the reason is clear when taking a look at the responses found on the corresponding CAIQ reports: CSP_3 replied either with a “YES” or a better quantitative security value much more questions than the other providers.

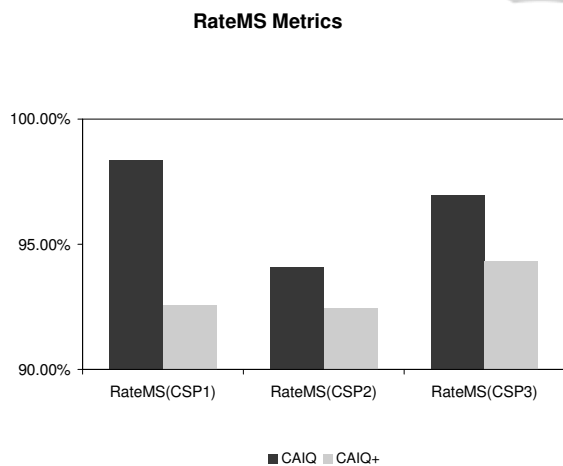


Figure 5: Using the RateMS metric to quantitatively compare CSPs. The higher RateMS is, the more security is being provided by the Cloud SecLA.

For our next test, let us suppose that a current customer of CSP_3 wants to know if either CSP_1 or CSP_2 might provide her with a better SecLA. For this purpose our hypothetical user first applies the metric proposed in Definition 2, using CSP_3 as a baseline. The obtained results (see Figure 6) show that with the two policy subsets (CAIQ and CAIQ+), both CSP_1 and CSP_2 had approximately 20% of their provisions with a higher security level than CSP_3 ’s. However, it is noticeable that the number of provisions with a lower se-

Comparing CSP1 and CSP2 versus CSP3

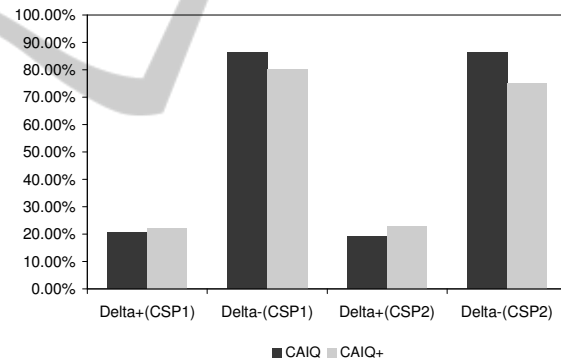


Figure 6: Quantitatively comparing CSPs with the metric from Definition 2.

Now the question that arises for the hypothetical customer of CSP_3 is: how the different CSPs compare on the individual categories of their respective SecLAs? This question can be answered through the *BenchScore* metric (cf. Definition 3), by *benchmarking* the Cloud SecLAs from both CSP_1 and CSP_2 using as a baseline CSP_3 . Obtained results are shown in Table 2, where the *BenchScore* metric was used for this purpose (column “Pass” in Table 2), but also was applied to quantitatively *benchmark* those Cloud SecLA provisions from CSP_1/CSP_2 which security level (i.e., LSL as explained in Section 3.1) fell behind the *benchmark* (column “Fail” in Table 2). The results shown in Table 2’s CAIQ+ column are consistent with our previous tests, because we observed that

Table 2: Benchmarking CSP_1 and CSP_2 versus CSP_3 .

	CSP_1				CSP_2			
	CAIQ		CAIQ+		CAIQ		CAIQ+	
	Pass	Fail	Pass	Fail	Pass	Fail	Pass	Fail
<i>BenchScore</i>	97.66%	2.34%	87.50%	12.50%	90.06%	9.94%	90.50%	9.50%
<i>CO</i>	7.02%	1.17%	7.00%	7.00%	7.60%	0.58%	12.50%	1.50%
<i>DG</i>	8.77%	0.00%	8.00%	0.00%	8.77%	0.00%	8.00%	0.00%
<i>FS</i>	5.26%	0.00%	4.50%	0.00%	4.68%	0.58%	4.00%	0.50%
<i>HR</i>	2.34%	0.00%	2.00%	0.00%	2.34%	0.00%	2.00%	0.00%
<i>IS</i>	41.52%	0.00%	38.00%	3.50%	37.43%	4.09%	38.00%	3.50%
<i>LG</i>	1.17%	0.00%	1.00%	0.00%	1.17%	0.00%	1.00%	0.00%
<i>OP</i>	2.92%	0.00%	2.50%	0.00%	2.92%	0.00%	2.50%	0.00%
<i>RI</i>	6.43%	0.58%	5.50%	0.50%	7.02%	0.00%	6.00%	0.00%
<i>RM</i>	2.92%	0.00%	2.50%	0.00%	1.75%	1.17%	1.50%	1.00%
<i>RS</i>	5.85%	0.58%	5.00%	1.00%	5.85%	0.58%	5.50%	0.50%
<i>SA</i>	13.45%	0.00%	11.50%	0.50%	10.53%	2.92%	9.50%	2.50%

the number of “failed” provisions increased for CSP_1 with respect to CSP_3 (from 2.34% — 12.50%).

As mentioned in Section 3 and seen on Figure 7, the proposed assessment metrics can be used at different levels of granularity varying from the Cloud SecLA-level to the individual security provision-level. With the quantitative level of detail being provided by the proposed metrics e.g., our hypothetical user should be able to take a decision about changing her current CSP depending on how important for her are the security provisions that “failed” the benchmarking process. This “weighting” of individual provisions is part of our future work, just as discussed in Section 6.

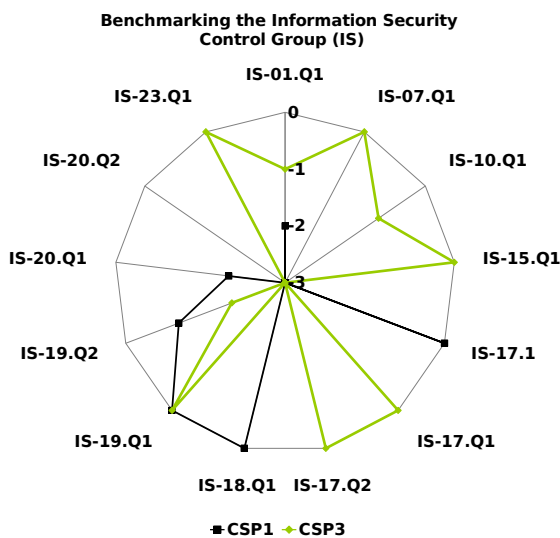


Figure 7: Benchmarking a subset of security provisions for two different CSPs.

5 RELATED WORK

SLA specifications and management is becoming an essential component within several emerging technologies such as Cloud computing. In the past, SLA specifications have been mainly considered in the Service Oriented Architectures (SOAs) and Web services fields like in (Andrieux, K., et.al.,) and (Ludwig, H., et.al.,). Unfortunately, one of the major limitations of these approaches is that they do not take into consideration security aspects even if the need for incorporating security in SLA was already highlighted some years ago (Henning, R., 1999).

Despite in the last years some proposals have started to consider security in SLAs, just very few of these have focused on Cloud SecLAs. In particular we highlight (Bernsmed, K., et.al., 2011), where the authors present a method for managing the SecLA lifecycle in the context of federated Cloud services. That work can be considered complementary to our research, taking into account that the authors of (Bernsmed, K., et.al., 2011) discuss the contents of Cloud SecLAs, but do not further elaborate about the techniques to quantitatively assess them.

In (Almorsy, M., et.al., 2011), the authors also propose the notion of evaluating Cloud SecLAs. In their paper, a metric is introduced to measure a CSP’s “security categorization” of their information (either per-tenant or per-service), based on impact metrics for all the three dimensions of security — confidentiality, integrity and availability —. However, the resulting security categorization is qualitative (i.e. specifying the High, Medium or Low security ranges), contrary to our quantitative security assessment metrics.

As mentioned in Section 1, to the best of our

knowledge there are no further works related with the quantitative security assessment of CSPs and in particular aimed to empirically validate their security metrics with real CSP data. Nevertheless for the sake of completeness, the rest of this section cites the efforts from other Information Technology fields (e.g. Web Services and Grid computing) aimed to adopt and assess SecLAs.

In (Frankova, G. and Yautsiukhin, A., 2007) and (Krautsevich, L., et.al., 2011), the authors propose a technique to aggregate security metrics from a web services' SecLA, however contrary to our research they did not propose the techniques to assess their SecLAs neither empirically validate the proposed metrics.

The Reference Evaluation Methodology — REM — (as explained in Section 3.1) was originally proposed in (Casola V., et.al., 2006). The authors introduced a methodology that can be adopted whenever there is the need of evaluating and comparing security SLAs (despite not specifically Cloud-related) expressed through the use of standard policy languages. A similar approach was used in (Casola V., et.al., 2005) and (Casola, V. et.al., 2007) to quantify the security of a Public Key Infrastructure, based on its Certificate Policy. The security assessment presented in our paper has been built above the methodology from (Casola V., et.al., 2006) by quantitatively evaluating a Cloud SecLA, but contrary to existing works we have also contributed with: (i) an additional set of security metrics to quantitatively assess the CSP, (ii) an initial Cloud SecLA specification and, (iii) the building blocks of an architecture aimed to empirically demonstrate the assessment of the CSP.

In (De Chaves, S. A., et.al., 2010) the authors highlight the importance of incorporating security metrics in SLAs and, in particular, of controlling and monitoring whether the security metrics are met. However, contrary to our paper, no further details are provided about the techniques used to represent and assess these SLAs.

A metric-based approach for assessing the security level of Critical Infrastructures was presented in (Ghani, H., et.al., 2010). In that article the authors define appropriate security metrics to monitor whether the established security requirements are fulfilled. Such metrics are also used for the definition of SLAs that should capture the defined requirements as well as the guarantees that the system provides together with the penalties that have to be applied when such guarantees are not met. The present paper took into account our previous experiences from (Luna, J., et.al., 2011), in order to contribute with the metrics and testbed to quantitatively assess Cloud SecLAs.

The works presented in (Irvine, C. and Levin, T., 2001) and (Neto, A., et.al., 2011) are also related with the assessment metrics presented in this paper. In (Irvine, C. and Levin, T., 2001), the authors support the notion of “Quality of Security” (similar to our SecLA quantification) and the usefulness of security ranges for quantifying security. Both arguments are directly related with the security quantification technique used by our framework (cf. Section 3). In (Neto, A., et.al., 2011) the authors present one of the few works related with security *benchmarking*. From their perspective *trust* can be used to *benchmark* security by performing the enumeration and accumulation of the evidence in favor of security that exists in a system, in the form of security mechanisms, processes, configurations, procedures and behaviors that prevent the accomplishment of specific threats that the system may be subjected to. These notions are directly related with the use of metrics for *benchmarking* security in our CSP assessment.

Finally, from an industrial perspective SecLAs must also be considered in security management regulations and standards as mentioned in (Monahan, B. and Yearworth, M., 2008). The Information Technology Infrastructure Library (ITIL) is one of the most prestigious good practices regarding IT Services Management and, some works have been done on modeling SLAs considering risks that threaten business processes and the requisites detailed in the Service Level Management module from ITIL (Feglar, T., 2004). However, contrary to the approach presented in our paper the techniques for quantitatively assessing these SecLAs in ITIL were not proposed in (Monahan, B. and Yearworth, M., 2008) nor (Feglar, T., 2004).

6 CONCLUSIONS

In this paper we have presented a practical CSP's security assessment use case based on (i) the concept of Security Level Agreements and, (ii) a novel set of quantitative security assessment metrics. This practical security assessment has been applied to the CSP's information stored in the STAR repository of the Cloud Security Alliance. To the best of our knowledge there are no previous works related with the quantitative security assessment of CSPs, and in particular aimed to empirically validate the applicability of their security metrics with real CSP data.

The obtained results show that the proposed security assessment can be performed at different levels of the CSP, ranging from the Cloud SecLA-level to the individual security provision-level. Our results

also demonstrate the need to create meaningful quantitative security metrics for CSPs, because despite the usefulness of initiatives like STAR, it is clear that the more information is provided by the CSP the more useful it will become to the end user. Fortunately, some ongoing works are taking place in work groups like CSA's Security Metrics WG (Cloud Security Alliance, 2011a) to create and motivate the use of these more quantitative security metrics for the Cloud.

We have adopted the REM technique to quantify Cloud SecLAs (as presented in Section 3.1), in particular due to the flexibility it offers to model and evaluate most classes of security policies. However, our future work will also consider the adoption of other SecLA evaluation techniques that might appear in the near future (e.g., the scoring model being developed by the Common Assurance Maturity Model community (Samani, R., et.al., 2011)).

The contributed security assessment metrics (cf. Section 3) are "Cloud SecLA-neutral", in the sense that they are able to work with different security criteria and standards. The only pre-requisite is to define the appropriate Cloud SecLA-templates to evaluate with the REM (as described in Section 3.1). Our current efforts in working groups like the CSA are aimed towards this objective. For example, we expect that in the near future it might be possible to create and assess SecLAs derived from both the threat analysis of real CSP architectures and, existing standards like e.g., ISO27001 and PCI.

Research works like the one presented in this paper will aid to trigger security transparency and new types of SecLA-based services in CSPs. Our belief is that SecLAs will become in the short term a key factor to enable competition between CSPs based on security properties.

The work in this paper is not aimed to substitute the Cloud security auditing function. In contrary, our security assessment rests on the belief that Cloud SecLAs are *trusted* in the sense that they have been previously audited by experts. This is precisely the base to build the *security assurance* needed by Cloud users, just as mentioned in Section 1.

The security assessment presented in this paper is part of a broader research focusing on *quantifying the end-to-end security in CSPs*, because the CSP's security level perceived by an end user will depend on the whole IT-chain, including the security behavior of the intermediate ISP, the requested CSP and the end user themselves. The most clear example is related to availability: despite a CSP advertises 99.99% of service availability, in the real-world it will depend on the availability of intermediate ISPs.

Future work will complement our security assess-

ment metrics (cf. Section 3) with the techniques to weight individual security provisions and to aggregate, negotiate, predict and tune Cloud SecLAs based not only on "declarative" information (e.g., the one from the STAR repository), but also on real-time data gathered from the service provider's infrastructure. The testbed presented in Section 4.1 can be further improved taking into account both this real-time feature and the idea of deploying publicly available security assessment services for the Cloud.

Our final goal is the creation of techniques and tools to empower end users through providing choices of service providers via the use of end-to-end security metrics. Users will be provided with adequate support to make informed decisions regarding the trustworthiness of an IT system to enable them to react depending on the changing conditions of the system in terms of security and trust. As a consequence, the confidence in the use of IT systems will increase. End users will also obtain a transparent view on the SecLAs agreed or expected from their service providers (Internet Service Providers included), thus re-balancing the current unequal relationship between both parties.

ACKNOWLEDGEMENTS

Research supported in part by EC FP7 IP ABC4TRUST and Loewe TUD CASED.

REFERENCES

- Almorsy, M., et.al. (2011). Collaboration-Based Cloud Computing Security Management Framework. In *Proc. of the IEEE International Conference on Cloud Computing*, pages 364–371.
- Andrieux, K., et.al. Web Services Agreement Specification (WS-Agreement). Technical Report TR-WSAgreement-2007, Open Grid Forum.
- Bernsmed, K., et.al. (2011). Security SLAs for Federated Cloud Services. In *Proc. of the IEEE Sixth International Conference on Availability, Reliability and Security*, pages 202–209.
- Boyle K., et.al. (2010). The CIS security metrics. Technical Report TR-28, Center for Internet Security.
- Casola V., et.al. (2005). A Reference Model for Security Level Evaluation: Policy and Fuzzy Techniques. *Journal of Universal Computer Science*, pages 150–174.
- Casola V., et.al. (2006). A SLA evaluation methodology in Service Oriented Architectures. In *Quality of Protection*, volume 23 of *Springer Advances in Information Security*, pages 119–130.
- Casola, V., et.al. (2007). Interoperable Grid PKIs Among Untrusted Domains: An Architectural Proposal. In

- Advances in Grid and Pervasive Computing*, volume 4459 of *Springer Lecture Notes in Computer Science*, pages 39–51.
- Casola, V. et.al. (2007). Static evaluation of Certificate Policies for Grid PKIs interoperability. In *Proc. of the IEEE Second International Conference on Availability, Reliability and Security*, pages 391–399.
- Center for Internet Security (2009). User Guide for CIS-CAT. Online: <http://benchmarks.cisecurity.org/en-us/docs/user-guides/CIS-CAT-Users-Guide.pdf>.
- Cloud Security Alliance (2011a). Security metrics workgroup. Online: <http://www.cloudsecurityalliance.org/Research.html>.
- Cloud Security Alliance (2011b). The Consensus Assessments Initiative Questionnaire. Online: <https://cloudsecurityalliance.org/research/cai/>.
- Cloud Security Alliance (2011c). The Security, Trust & Assurance Registry (STAR). Online: <https://cloudsecurityalliance.org/star/>.
- De Chaves, S. A., et.al. (2010). SLA perspective in security management for Cloud computing. In *Proc. of the IEEE Sixth International Conference on Networking and Services*, pages 212–217.
- Dekker, M. and Hogben, G. (2011). Survey and analysis of security parameters in cloud SLAs across the European public sector. Technical Report TR-2011-12-19, European Network and Information Security Agency.
- Feglar, T. (2004). ITIL based Service Level Management if SLAs cover Security. *Journal on Systemics, Cybernetics and Informatic*, pages 61–71.
- Frankova, G. and Yautsiukhin, A. (2007). Service and protection level agreements for business processes. In *Proc. of the IEEE Second European Young Researchers Workshop on Service Oriented Computing*, page 38.
- Ghani, H., et.al. (2010). Assessing the Security of Internet Connected Critical Infrastructures (The CoMiFin Project Approach). In *Proc. of the Workshop on Security of the Internet of Things*.
- Henning, R. (1999). Security service level agreements: quantifiable security for the enterprise? In *Proc. of the ACM Workshop on New security paradigms*, pages 54–60.
- Irvine, C. and Levin, T. (2001). Quality of security service. In *Proc. of the ACM Workshop on New security paradigms*, pages 91–99.
- Jansen, W. (2010). Directions in security metrics research. Technical Report TR-7564, National Institute for Standards and Technology.
- Krautsevich, L., et.al. (2011). A general method for assessment of security in complex services. In *Towards a Service-Based Internet*, volume 6994 of *Springer Lecture Notes in Computer Science*, pages 153–164.
- Ludwig, H., et.al. Web Service Level Agreement (WSLA) Language Specification. Technical Report TR-WSLA-2003-01-28, IBM.
- Luna, J. et.al. (2008). Providing security to the Desktop Data Grid. In *Proc. of the IEEE International Symposium on Parallel and Distributed Processing*, pages 1–8.
- Luna, J., et.al. (2011). A Security Metrics Framework for the Cloud. In *Proc. of the INSTICC International Conference on Security and Cryptography*, pages 245–250.
- Monahan, B. and Yearworth, M. (2008). Meaningful security SLAs. Technical Report TR-HPL-2005-218, HP Labs.
- Neto, A., et.al. (2011). To benchmark or not to benchmark security: That is the question. In *Proc. of the IEEE Dependable Systems and Networks Workshops*, pages 182–187.
- Samani, R., et.al. (2011). Common Assurance Maturity Model: Scoring Model. Online: <http://common-assurance.com/>.
- Trimintzios, P. (2011). Measurement Frameworks and Metrics for Resilient Networks and Services. Discussion Draft. European Network and Information Security Agency.
- Verendel, V. (2009). Quantified security is a weak hypothesis: a critical survey of results and assumptions. In *Proc. of the ACM Workshop on new security paradigms*, pages 37–50.
- Weisstein, W. (2011a). Frobenius Norm. Online: <http://mathworld.wolfram.com/FrobeniusNorm.html>.
- Weisstein, W. (2011b). L1-Norm. Online: <http://mathworld.wolfram.com/L1-Norm.html>.