

RESEARCH

Open Access



A risk assessment model for selecting cloud service providers

Erdal Cayirci¹, Alexandr Garaga^{2,3}, Anderson Santana de Oliveira^{2*} and Yves Roudier³

Abstract

The Cloud Adoption Risk Assessment Model is designed to help cloud customers in assessing the risks that they face by selecting a specific cloud service provider. It evaluates background information obtained from cloud customers and cloud service providers to analyze various risk scenarios. This facilitates decision making in selecting the cloud service provider with the most preferable risk profile based on aggregated risks to security, privacy, and service delivery. Based on this model we developed a prototype using machine learning to automatically analyze the risks of representative cloud service providers from the Cloud Security Alliance Security, Trust & Assurance Registry.

Keywords: Risk assessment, Cloud computing, Security, Privacy

Abbreviations: A4Cloud, Accountability for cloud and other future internet services; CAIQ, Consensus assessment initiative questionnaire; CARAM, Cloud adoption risk assessment model; CCM, Cloud control matrix; CNIL, Commission nationale de l'informatique et des libertés; CSA, Cloud security alliance; CSC, Cloud service consumer; CSP, Cloud service provider; ENISA, European network and information security agency; IEC, International electrotechnical commission; ISACA, Information systems audit and control association; ISO, International organization for standardization; IT, Information technology; JRTM, Joint risk and trust model; NIST, National institute of standards and technology; SeCLA, Cloud security level agreements; SMB, Small-medium business; STAR, Security, trust & assurance registry

Introduction

Moving business processes to the cloud is associated with a change in the risk landscape to an organization [1]. Cloud Security Alliance (CSA) [2] has found that insufficient due diligence was among the top threats in cloud computing in 2013. This threat is linked to the fact that organizations which strive to adopt cloud computing often do not understand well the resulting risks.

Regulations related to data protection, financial reporting, etc. put certain requirements that should be complied with even when outsourcing business processes to 3rd parties, like cloud service providers (CSPs). For example, EU Data Protection Directive, in particular Article 29 Data Protection Working Party [3] recommends that all data controllers (usually corporate cloud customers) perform

an impact assessment of moving personal data of their clients to the cloud.

However, most of the cloud service customers (CSCs), especially Small-Medium Businesses (SMBs), may not have enough knowledge in performing such assessments at a good level, because they may not necessarily employ IT specialists and the lack of transparency is intrinsic to the operations of the CSPs. This makes difficult to choose an appropriate CSP based on CSC's security requirements, especially considering the abundance of similar cloud offerings [4].

This work proposes a methodology, cloud adoption risk assessment model (CARAM), to help in assessing the various risks to business, security and privacy that CSCs face when moving to the cloud by leveraging information from CSCs, CSPs and several public sources. CARAM consists of the following tools that complement the various recommendations from European Network and Information Security Agency (ENISA) [1], and Cloud

*Correspondence: anderson.santana.de.oliveira@sap.com

²SAP Labs France, Mougins, France

Full list of author information is available at the end of the article

Security Alliance (CSA) for a complete risk assessment framework:

- A questionnaire for CSCs
- A tool and an algorithm to classify the answers to Cloud Assessment Initiative Questionnaire (CAIQ) to discrete values
- A model that maps the answers to both questionnaires to risk values
- A multi-criteria decision approach with posterior articulation of CSC preferences for relative risk analysis, using a few parameters for security, privacy and quality of service, allowing to quickly and reliably compare multiple CSPs

This paper extends our work in [5] with experimental results - we devised profiles representing realistic customer categories to classify providers according to the CSC needs. We also used a more precise risk scale for comparing CSPs, allowing one to visualize the differences in security practices of the most representative players in the cloud services landscape. Therefore the current version brings significant improvements with respect to the previous paper [5].

In Section “Related work” we elaborate on the literature related to the risk assessment for adoption of cloud computing: we focus on the work carried out by ENISA and CSA because CARAM is based on them; In Section “Risk levels computation” we introduce CARAM, and then a multi-criteria risk assessment approach with posterior articulation of the CSCs; In Section “Experimental results” we demonstrate experimental results from using CARAM on a case study; In Section “Limitations” we outline some limitations of the approach; We conclude our paper in Section “Conclusion”.

Related work

Several large standardization bodies such as International Organization for Standardization (ISO), International Electrotechnical Commission (IEC) and National Institute of Standards and Technology (NIST) and Information Technology (IT) Governance Institute and the Information Systems Audit and Control Association (ISACA) published standards on IT risk management and risk assessment: ISO 31000 [6], ISO/IEC 31010 [7], ISO/IEC 27005 [8], NIST SP 800-30 [9], SP 800-37 [10] and COBIT [11]. All these standards are generic i.e. not specific to cloud deployments, and while possible to use them for evaluating different cloud solutions, it will require a considerable amount of effort and expert knowledge, which SMBs cannot always afford.

Some adaptations of these standards were developed specifically for cloud deployments. E.g. Microsoft proposed a Cloud Decision Framework [12] based on ISO 31000. It provides guidance for risk assessment to be

performed by potential CSCs when choosing a cloud solution. The risk profiles of different cloud solutions are constructed based on a predefined set of risks from four categories: compliance, strategic, operational and market and finance. The authors suggest using CSA Cloud Control Matrix (CCM) [13] as guidance for evaluating mitigating controls. While this approach could be more practical than other generic risk assessment frameworks for evaluating different cloud solutions it is still quite abstract and largely rely on experts’ opinions for estimating and evaluating the risks and mitigations. In contrast, we propose a concrete step by step approach to automate the estimation and evaluation of risks of adopting different cloud solutions. This could be more suitable for smaller organizations that lack sufficient resources for a full-scale risk assessment.

ENISA [1] provided recommendations and a framework for generic qualitative inductive risk assessment for cloud computing. Their recommendations include extensive lists of possible incident scenarios, assets and vulnerabilities in cloud computing deployments. It suggests estimating risk levels on the basis of likelihood of a risk scenario mapped against the estimated negative impact, which is the essence of the risk formulation by also many others in the literature [7, 11, 14, 15]. Although ENISA’s recommendations are specific for cloud computing, it is a generic framework that does not provide an approach to map the specifics of CSPs and CSCs to the 35 risk scenarios listed in the report [1]. In Section “Risk levels computation” we describe how CARAM fine-tunes this approach to estimate risk values based on some known information about CSCs and CSPs.

Another qualitative inductive scheme was published by “The Commission nationale de l’informatique et des libertés” (CNIL) or in English: The French National Commission on Informatics and Liberty [16] more recently. CNIL’s methodology is similar to the ENISA’s framework with the following difference: it is a risk assessment focused on privacy risks in cloud computing. It also recommends measures to reduce the risks and assess the residual privacy risks after the application of these measures. However, it is still generic and does not account for specific requirements of CSPs or CSCs.

The CSA Cloud Assessment Initiative Questionnaire (CAIQ) [17] is a questionnaire prepared for CSPs to document the implemented security measures. It is based on the CSA Cloud Control Matrix (CCM) taxonomy of security controls [13] and is aimed to help CSCs understand the security coverage of specific cloud offerings in relation to popular security standards, control frameworks and regulations. The questionnaires answered by many CSPs are publicly available in the CSA Security, Trust and Assurance Registry (STAR) [4]. We propose a methodology that uses the data extracted from STAR to

evaluate the implementation of various controls provided by cloud solutions (see Subsection “The vulnerability parameter for a CSP”).

Luna et al. introduced in [18] Cloud Security Level Agreements (SecLA) and proposed a methodology to benchmark SecLA of CSPs with respect to CSCs’ requirements [19]. Both CSP SecLA provisions and user requirements are expressed using a special data structure: Quantitative Policy Trees, allowing expressing controls with different granularity: CCM control areas, control groups, and controls (corresponding to CAIQ answers). The authors demonstrate their approach using data on several CSPs from STAR, by calculating security levels for respective controls and control groups. While similar in the intent CARAM is a model for risk assessment, while [19] proposes a ranking algorithm for matching CSC requirements vs. CSP provisions. In [19] CSCs need a certain level of security expertise to specify their requirements, while in CARAM this is not necessary: CSCs only need to specify acceptable risk levels for security, privacy and service categories, while still allowing a more fine grained specification. Another major difference is that [19] assumes the existence of a mapping from provisions to quantitative Local Security Levels to allow further analysis. Given a high number of potential CSPs and controls for each CSP creating this mapping would require significant manual work. In CARAM we propose a way to automatically construct such a mapping (see Subsection “The vulnerability parameter for a CSP”).

Habib et al. proposed a multi-faceted Trust Management system architecture for a cloud computing marketplace [20]. The system evaluates the trustworthiness of CSPs in terms of different SLA attributes assessed using information collected from multiple sources. This is done by evaluating opinions related to SLA attributes and aggregating them into a trust score for a CSP. The authors mention CAIQ answers as a source of information, however they do not specify how exactly the CSP trust score is computed from the answers, especially considering that the answers are in free text form.

In [21], the EU funded project SECCRIT enumerated very relevant cloud risk scenarios systematically, in a similar fashion to what ENISA [1] did a few years earlier, but this time they evaluated the risk perception from the users point of view. They used a survey-based methodology to ask respondents to rank risks in a standard way - assessing probability and impact. CARAM is innovative when compared with this approach in two ways: it is using real information disclosed by cloud providers to estimate the likelihood of threats affecting the cloud would become concrete. Second, it allows the user of the methodology to focus on what they know best, their assets: the methodology only requires to assign priorities to assets in order to provide quantifiable risk information.

Other EU projects such as SPECS [22] and Escudo-Cloud [23] also looked at quantifying the capacity of a given cloud provider to satisfy security Service Level Agreements (SLAs) using information published in the CSA STAR registry. It also has a focus on usability, but the endeavour is different in essence. Here we are assessing risk scenarios in an automated way, whereas in that work [23] does not provide an automated processing for the answers given by the CSPs to the CAIQ, making it difficult to compare more than three providers quickly.

Joint Risk and Trust Model (JRTM) [24] was developed by Accountability for Cloud and Other Future Internet Services Project (A4Cloud). JRTM is a quantitative risk assessment model that assesses the cloud service security and privacy risks for a specific CSP and CSC. It counts on a third party (i.e., a Trust as a Service Provider) to accumulate statistical data (i.e., evidence) on the trustworthiness of CSPs. These evidences include the number of security, privacy and service events that a CSP was subject to and the percentage of the events that the CSP recovered from before they become an incident (i.e., they impact on CSC). However, such detailed statistical data is not always available and even if available they will not be normally shared with (potential) CSCs—hence the need for a trusted third party. In this work, instead, we rely on public information already provided by CSPs regarding their implementation of various security controls for a qualitative risk assessment of their solutions. This can be considered as an extension of JRTM when statistical data is available or a substitution otherwise.

One of the main advantages of CARAM is that it is easy to make it evolve to cover further risk scenarios and to make it evolve if the security control frameworks evolve. As we will introduce in the next section, the risk level computation is parametrized by mappings between threats and security controls that help to mitigate them, making the approach relevant over time.

Risk levels computation

ENISA [1] identified 35 incident scenarios that fall in one of the following four categories: policy and organizational, technical, legal and the other scenarios not specific to cloud computing (see Table 1). The likelihood of each of these scenarios and their business impact are determined in consultation with an expert group. The scale of probability and impact has five discrete classes between very low and very high. For example, the probability and impact of Incident Scenario P1 in “Policy and Organizational Scenarios” category (i.e., lock-in) are given as HIGH and MEDIUM relatively.

ENISA also provides a list of 53 vulnerabilities (i.e., 31 cloud specific and 22 not cloud specific vulnerabilities, see Table 2) and 23 classes of CSC assets (see Table 3) that may be affected by the cloud adoption. Each of 35 incident

Table 1 ENISA's list of risk scenarios and their categories

Risk category	Risk name
Policy & Organizational	P1. Lock-in
	P2. Loss of governance
	P3. Compliance challenges
	P4. Loss of business reputation due to co-tenant activities
	P5. Cloud service termination or failure
	P6. Cloud provider acquisition
	P7. Supply chain failure
Technical	T1. Resource exhaustion (under or over provisioning)
	T2. Isolation failure
	T3. Cloud provider malicious insider - abuse of high privilege roles
	T4. Management interface compromise (manipulation, availability of infrastructure)
	T5. Intercepting data in transit
	T6. Data leakage on up/download, intra-cloud
	T7. Insecure or ineffective deletion of data
	T8. Distributed denial of service (DDoS)
	T9. Economic denial of service (EDoS)
	T10. Loss of encryption keys
	T11. Undertaking malicious probes or scans
	T12. Compromise service engine
	T13. Conflicts between customer hardening procedures and cloud environment
Legal	L1. Subpoena and e-discovery
	L2. Risk from changes of jurisdiction
	L3. Data protection risks
	L4. Licensing risks
Not Specific to the Cloud	N1. Network breaks
	N2. Network management (ie, network congestion / mis-connection / non-optimal use)
	N3. Modifying network traffic
	N4. Privilege escalation
	N5. Social engineering attacks (ie, impersonation)
	N6. Loss or compromise of operational logs
	N7. Loss or compromise of security logs (manipulation of forensic investigation)
	N8. Backups lost, stolen
	N9. Unauthorized access to premises (including physical access to machines and other facilities)
	N10. Theft of computer equipment
	N11. Natural disasters

Table 2 ENISA's list of vulnerabilities

Cloud specific vulnerabilities
V1. Authentication Authorization Accounting (AAA) vulnerabilities
V2. User provisioning vulnerabilities
V3. User de-provisioning vulnerabilities
V4. Remote access to management interface
V5. Hypervisor vulnerabilities
V6. Lack of resource isolation
V7. Lack of reputational isolation
V8. Communication encryption vulnerabilities
V9. Lack of or weak encryption of archives and data in transit
V10. Impossibility of processing data in encrypted form
V11. Poor key management procedures
V12. Key generation: low entropy for random number generation
V13. Lack of standard technologies and solutions
V14. No source escrow agreement
V15. Inaccurate modelling of resource
V16. No control on vulnerability assessment process
V17. Possibility that internal (cloud) network probing will occur
V18. Possibility that co-residence checks will be performed
V19. Lack of forensic readiness
V20. Sensitive media sanitization
V21. Synchronizing responsibilities or contractual obligations external to cloud
V22. Cross-cloud applications creating hidden dependency
V23. SLA clauses with conflicting promises to different stakeholders
V24. SLA clauses containing excessive business risk
V25. Audit or certification not available to customers
V26. Certification schemes not adapted to cloud infrastructures
V27. Inadequate resource provisioning and investments in infrastructure
V28. No policies for resource capping
V29. Storage of data in multiple jurisdictions and lack of transparency about this
V30. Lack of information on jurisdictions
V31. Lack of completeness and transparency in terms of use
Vulnerabilities not specific to the cloud
V32. Lack of security awareness
V33. Lack of vetting processes
V34. Unclear roles and responsibilities
V35. Poor enforcement of role definitions
V36. Need-to-know principle not applied
V37. Inadequate physical security procedures
V38. Misconfiguration
V39. System or OS vulnerabilities
V40. Untrusted software

Table 2 ENISA's list of vulnerabilities (Continued)

V41. Lack of, or a poor and untested, business continuity and disaster recovery plan
V42. Lack of, or incomplete or inaccurate, asset inventory
V43. Lack of, or poor or inadequate, asset classification
V44. Unclear asset ownership
V45. Poor identification of project requirements
V46. Poor provider selection
V47. Lack of supplier redundancy
V48. Application vulnerabilities or poor patch management
V49. Resource consumption vulnerabilities
V50. Breach of NDA by provider
V51. Liability from data loss
V52. Lack of policy or poor procedures for logs collection and retention
V53. Inadequate or misconfigured filtering resources

scenarios is related with a subset of vulnerabilities and assets. For example, the Incident Scenario P1 is related to Vulnerabilities V13 (lack of standard technologies and solutions), V31 (lack of completeness and transparency in terms of use), V46 (poor provider selection), V47 (lack

Table 3 ENISA's list of assets

A1. Company reputation
A2. Customer trust
A3. Employee loyalty and experience
A4. Intellectual property
A5. Personal sensitive data
A6. Personal data
A7. Personal data - critical
A8. HR data
A9. Service delivery - real time services
A10. Service delivery
A11. Access control / authentication / authorization (root/admin v others)
A12. Credentials
A13. User directory (data)
A14. Cloud service management interface
A15. Management interface APIs
A16. Network (connections, etc.)
A17. Physical hardware
A18. Physical buildings
A19. Cloud Provider Application (source code)
A20. Certification
A21. Operational logs (customer and cloud provider)
A22. Security logs
A23. Backup or archive data

of supplier redundancy) and Assets A1 (company reputation), A5 (personal sensitive data), A6 (personal data), A7 (personal data critical), A9 (service delivery - real time services), A10 (service delivery).

The likelihood and business impact values that are determined by the experts are converted to the risk levels for each incident scenario based on a risk matrix with a scale between 0 and 8 as shown in Fig. 1. Then, the risk levels are mapped to a qualitative scale as follows:


- Low → 0-2
- Medium → 3-5
- High → 6-8

Hence, a CSC can assess the risk level related to an incident scenario qualitatively and understands what kind of vulnerabilities and assets are related to each scenario by examining [1]. These values represent educated guesses over a wide range of common cloud deployments and do not have a precise semantics. In practice, the risk levels are related to many factors such as the security controls that CSPs implement and the concerned assets of the specific users. Therefore, a generic value cannot be applied to all CSPs and CSCs. Although vulnerabilities and assets for each incident scenario are given by the ENISA framework, it does not describe how those values can be adapted for a specific CSP and CSC pair. CARAM fills this gap. For that, first the qualitative scale used by ENISA as probability and impact values are mapped to a quantitative scale as follows:

- Very low → 1
- Low → 2
- Medium → 3
- High → 4
- Very high → 5

For example, probability P_1 and impact I_1 values for the first scenario (i.e., lock in) is HIGH and MEDIUM respectively. We map these values as follows: $P_1 = 4$ and $I_1 = 3$.

To compute the risk levels considering particular vulnerabilities and assets involved CARAM adjusts the values from ENISA, taken as a baseline, taking into account



additional information about the cloud service. For that, we use Eqs. 1 and 2:

$$\beta_i = P_i \times v_i \tag{1}$$

$$\delta_i = I_i \times \alpha_i \tag{2}$$

In Eq. 1, for the risk scenario i , β_i is the adjusted probability, v_i is the vulnerability index of a given CSP, δ_i is the adjusted impact and α_i is the asset index for a given CSC. Here we assume that probability and impact of an incident are proportional to the number of non-addressed vulnerabilities by a CSP; and impact is proportional to the number of CSC assets related to a risk scenario.

Note that vulnerability index of a CSP is the same for all CSCs and the asset index of a CSC is the same for all CSPs. Vulnerability and asset indices are calculated as given in Eqs. 3 and 4 respectively, where $v_{ki} = 1$ if vulnerability k is in the list of vulnerabilities [1] for risk scenario i , and 0 otherwise. Similarly, $\alpha_{ki} = 1$ if asset k is in the list of assets [1] for risk scenario i . Please note again that there are 53 vulnerabilities (see Table 2) and 23 assets (see Table 3) listed in [1]. The vulnerability related parameter ϵ_k in Eq. 3 is derived from the answers to CAIQ, and elaborated on later in Subsection “The vulnerability parameter for a CSP”. The asset related parameter γ_k in Eq. 4 is given value 0 if the CSC’s answer to the question “Does the service that you seek will involve any asset of yours that fall in the same category as asset k ?” is “No”, and value 1 otherwise.

$$v_i = \frac{\sum_{k=1}^{53} v_{ki} \times \epsilon_k}{\sum_{k=1}^{53} v_{ki}} \tag{3}$$

$$\alpha_i = \frac{\sum_{k=1}^{53} \alpha_{ki} \times \gamma_k}{\sum_{k=1}^{53} \alpha_{ki}} \tag{4}$$

We would like to highlight that CARAM is independent from the number of incident scenarios and probability, impact, vulnerability and assets assigned to the incident scenarios. Moreover, it is possible to assign weight values for each of assets and vulnerabilities if some of them are assumed as of higher importance comparing to the others.

The vulnerability parameter for a CSP

We use CSPs’ responses to CAIQ from [4] to assign a value to the vulnerability related parameter ϵ_k . Most of the entries in STAR are using the CAIQ v1.1 template [25] that provides 148 questions grouped into the control areas shown in Table 4 covering the state of implementation of security controls.

The v1.1 template accepts a free text answer in contrast to the newer v3.0.1 template where the CSPs are expected to choose their answers between “Yes”, “No” and

Table 4 The control areas in CAIQ

Compliance (CO)	Operations Management (OP)
Data Governance (DG)	Risk Management (RI)
Facility Security (FS)	Release Management (RM)
Human Resources Security (HR)	Resiliency (RS)
Information Security (IS)	Security Architecture (SA)
Legal (LG)	

“Not applicable”. This makes the answers unsuitable for automated analysis. We proposed a mechanism to map the answers given to the questions in CAIQ to one of the control implementation status categories in Table 5.

The category *Implemented* has a positive meaning (the control is in place), but the answer “Yes” to a CAIQ question does not always imply a more secure system. For example, the “Yes” answer to CAIQ Question RS06-01 “Are any of your data centers located in places which have a high probability/occurrence of high-impact environmental risks (floods, tornadoes, earthquakes, hurricanes, etc.)?” implies a negative outcome, which means the control is not implemented. In such cases we classify the “Yes” answer as *Not Implemented* and the “No” answer as *Implemented*.

We analyzed 44 out of 70 CSPs from the mentioned registry providing answers to about 200 questions each. To be included in the selection a provider had to fill in the CAIQ questionnaire in the format provided by CSA (to enable automated processing). The responses of some of the big CSPs (e.g. Amazon, HP, Microsoft, RedHat, or SAP¹) who provided answers in other forms, though, were processed manually to ensure the consideration of the major players. Given the workload, we decided to automate the classification of the free text answers to CAIQ questions using the supervised machine learning algorithms (sequential minimal optimization and string vectorization) provided by the WEKA tool [26]. We created a training set from a random sampling of around 300 manually classified answers out of overall circa 9000 answers and used it to classify the other remaining answers. The 10-fold cross-validation provided an accuracy of around 84 % of correctly classified instances, which we consider enough for our purpose.

Table 5 The categorization of the answers given to the questions in CAIQ

<i>Implemented</i>	the control is in place
<i>Conditionally implemented</i>	the control can be implemented under some conditions
<i>Not implemented</i>	the control is not in place
<i>Not applicable</i>	the control is not applicable to the provided service

After classification of the answers to one of the categories in Table 5, the implementation value q_m is assigned for each of the controls identified by question m :

- *Implemented* $\rightarrow q_m = 0$: the related vulnerabilities are mitigated
- *Not Applicable* $\rightarrow q_m = 0$: these controls do not impact the risk value
- *No* $\rightarrow q_m = 1$: the related vulnerabilities are not mitigated
- *Conditionally Implemented* \rightarrow the CSC needs to clarify with the CSP if the control can be implemented. If yes, $q_m = 0$. Otherwise, $q_m = 1$.

When q_m is known for a CSP and a CSC, Eq. 5 gives the vulnerability related parameter ϵ_k for the CSP and the CSC. Please note that this value is for a specific CSP and CSC pair.

$$\epsilon_k = \frac{\sum_{m=1}^n r_{mk} \times q_m}{\sum_{m=1}^n r_{mk} \times b_m} \quad (5)$$

In Eq. 5, n is the number of questions in CAIQ. r_{mk} is the mapping of the CAIQ questions to vulnerabilities: it is 1 if the control representing the question m mitigates the vulnerability k , and 0 otherwise (see Table 6 for an example mapping excerpt and Additional file 1 — Mapping CAIQ questions to vulnerabilities.xlsx for the full mapping).

Finally, $b_m = 0$ if the answer to the question m is “Not Applicable” and 1 otherwise. This allows discarding the unrelated questions to avoid wrongly penalizing the CSPs.

In Eq. 5 ϵ_k receives a minimum value 0 if all the controls related to the vulnerability k are implemented and hence the vulnerability does not impact negatively the risk

values. The more controls related to the vulnerability k are not implemented, the higher ϵ_k is. Its maximum value is 1, which means the CSP has no measures against the vulnerability k .

Relative risk assessment-based CSP selection

ENISA’s risk assessment model is based on 35 incident scenarios. These numerous risks make it difficult for customers to meaningfully compare multiple providers. Therefore, we first reduce the number of criteria from these 35 incident scenarios to three categories of cloud risks: security, privacy and service [24]. For that, we compute the probability that a privacy (β_r), a security (β_s) and a service (β_e) incident can occur and the impact of a privacy (δ_r), a security (δ_s) and a service (δ_e) incident by applying Eqs. 6 to 11. In Eqs. 6 and 9, $r_i = 1$ if ENISA incident scenario i is related to privacy, and 0 otherwise (see Table 7 for an example mapping). ω_{ri} and α_{ri} are real numbers between 0 and 1. They are the weight factors for probability and impact respectively. The significance of every scenario may not be the same when calculating an aggregated value for privacy, security and service incidents. Moreover, the scenarios may need to be treated differently for each CSC especially when calculating the aggregated impact values. The weight factors are for making these adjustments. If the significance of each scenario is the same, then the weight factors can be assigned 1. Similar to r_i , s_i and e_i are the mapping values for security and service risks respectively (see Table 7). ω_{si} and α_{si} are the weight factors for security scenarios, and ω_{ei} and α_{ei} are the weight factors for service scenarios.

$$\beta_r = \frac{\sum_{i=1}^{35} \beta_i \times r_i \times \omega_{ri}}{\sum_{i=1}^{35} r_i \times \omega_{ri}} \quad (6)$$

$$\beta_s = \frac{\sum_{i=1}^{35} \beta_i \times s_i \times \omega_{si}}{\sum_{i=1}^{35} s_i \times \omega_{si}} \quad (7)$$

$$\beta_e = \frac{\sum_{i=1}^{35} \beta_i \times e_i \times \omega_{ei}}{\sum_{i=1}^{35} e_i \times \omega_{ei}} \quad (8)$$

$$\delta_r = \frac{\sum_{i=1}^{35} \delta_i \times r_i \times \alpha_{ri}}{\sum_{i=1}^{35} r_i \times \alpha_{ri}} \quad (9)$$

$$\delta_s = \frac{\sum_{i=1}^{35} \delta_i \times s_i \times \alpha_{si}}{\sum_{i=1}^{35} s_i \times \alpha_{si}} \quad (10)$$

$$\delta_e = \frac{\sum_{i=1}^{35} \delta_i \times e_i \times \alpha_{ei}}{\sum_{i=1}^{35} e_i \times \alpha_{ei}} \quad (11)$$

Table 6 Mapping r_{mk} of CAIQ questions to ENISA vulnerabilities (excerpt)

Control group	Vulnerabilities mitigated
Audit Planning CO-01	V02, V03, V13, V14, V16, V23, V25, V26, V27, V29, V33, V35, V50
Independent Audits CO-02	V02, V03, V13, V14, V16, V23, V25, V26, V27, V29, V33, V35, V50
Third Party Audits CO-03	V02, V03, V13, V14, V16, V23, V25, V26, V27, V29, V33, V35, V50
Contact/Authority Maintenance CO-04	V14, V21, V29, V30
Information System Regulatory Mapping CO-05	V07, V08, V09, V10
Intellectual Property CO-06	V34, V31, V35, V44
Intellectual Property CO-07	V34, V31, V35, V44
Intellectual Property CO-08	V34, V31, V35, V44

Table 7 Mapping (r_i, s_i, e_i) of ENISA risk scenarios to risk categories

Risk scenario i	Privacy r_i	Security s_i	Service e_i
P1	0	0	1
P2	1	0	0
P3	1	1	1
P4	0	1	0
P5	0	0	1
P6	1	1	1
P7	0	0	1
T1	0	0	1
T2	1	1	0
T3	1	1	1
T4	1	1	1
T5	1	1	0
T6	1	1	0
T7	1	1	0
T8	0	0	1
T9	0	0	1
T10	1	1	0
T11	1	1	0
T12	1	1	1
T13	0	1	0
L1	1	1	0
L2	1	0	0
L3	1	1	0
L4	0	0	1
N1	0	0	1
N2	0	0	1
N3	0	0	1
N4	1	1	1
N5	0	1	0
N6	0	1	1
N7	0	1	1
N8	1	1	1
N9	1	1	0
N10	1	1	1
N11	0	0	1

When probability (i.e., β) and impact (i.e., δ) values are calculated, they are mapped to the qualitative scale as follows:

- [0, 0.5) → Negligible
- [0.5, 1) → Extremely Low
- [1, 1.5) → Very Low
- [1.5, 2) → Low
- [2, 2.5) → Below Average

- [2.5, 3) → Above Average
- [3, 3.5) → High
- [3.5, 4) → Very High
- [4, 4.5) → Extremely High
- [4.5, 5] → Not Recommended

Please note that this is a higher resolution scale with 10 values comparing to ENISA's original five value qualitative scale. We need a higher resolution scale to differentiate CSPs because the adjusted probabilities of risk scenarios are mostly below the average for the CSPs that answer CAIQ. Those CSPs are clearly aware of the incident scenarios and implement at least a subset of the controls, which are subject in CAIQ. Finally, by using the matrix in Fig. 2—which is similar to but has a higher resolution than the matrix in Fig. 1—the risk values for privacy R_r , security R_s and service R_e are obtained. Please note that these values are calculated for each CSP-CSC pair. Although, the color codes in Fig. 2 are only for the three value qualitative scale (i.e., Low < Medium < High) as in the ENISA's assessment, the quantitative scale for the overall risk assessment is also higher resolution (between 0 and 18 instead of 0 and 10) which is the result of selecting higher resolution scales for probability and impact.

At this stage, the CSC provides CARAM with the maximum acceptable levels of risks for privacy R_{rmax} , security R_{smax} and service R_{smax} . The CSC may also provide a set $U = \{p_1, \dots, p_n\}$ of CSPs that should be excluded from the assessment due to reasons like business relations, politics, past experience, etc. When this information is available, CARAM creates a set F of feasible CSPs out of the set S of all the CSPs available for assessment (i.e. CSPs that have a completed CAIQ in STAR) such that $F \subset S$ using Eq. 12.

$$p_i \in F \iff (p_i \notin U) \wedge (R_{rmax} > R_{ri}) \wedge (R_{smax} > R_{si}) \wedge (R_{emax} > R_{ei}) \tag{12}$$

where R_{ri} , R_{si} and R_{ei} are the privacy, security and service risks for the CSP p_i .

F can be an empty set, a set with only one element or multiple elements. If F is an empty set, there is no feasible solution for the CSC. If F has only one element, that is the only feasible solution for the CSC under the given constraints. In both of these cases, CARAM informs the CSC directly with the result. If F has multiple elements, all the dominating CSPs are removed from F resulting in the set F' . Here we define the dominating relation $>$: $CSP_1 > CSP_2 \iff R_{ri}, R_{si}$ and R_{ei} for CSP_1 are higher than those for CSP_2 . If the resulting F' includes only one CSP, CARAM informs the CSC about the solution that fits the best to it. If there are multiple CSPs in F' , the CSC is given the complete F' for the posterior articulation of the preferences.

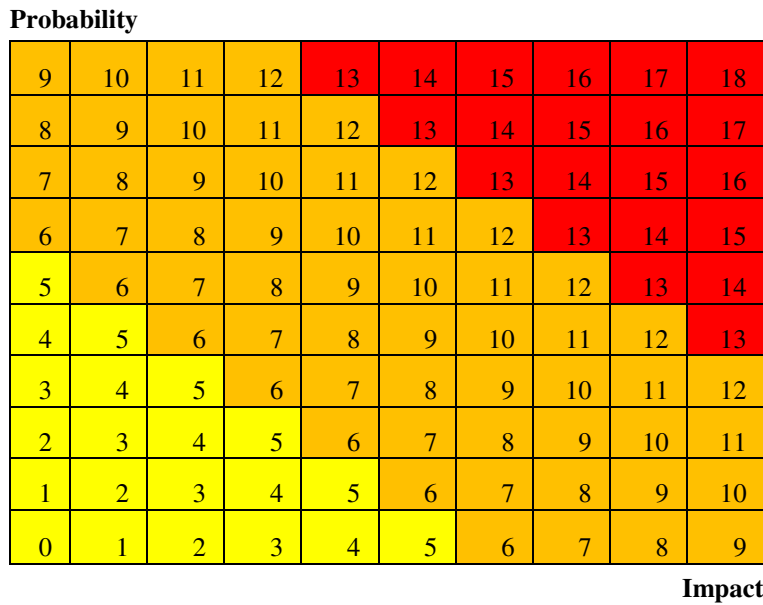


Fig. 2 CARAM definition of risk levels

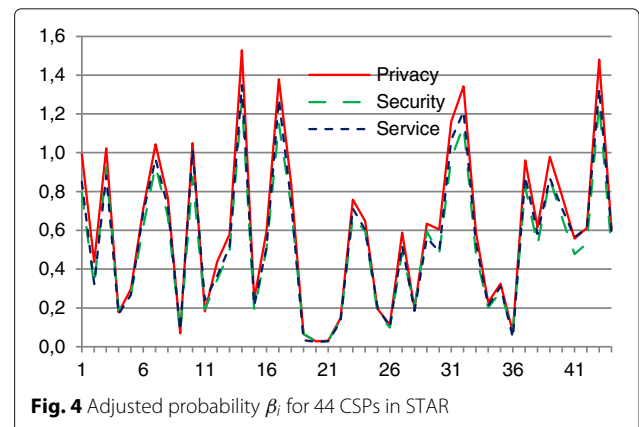
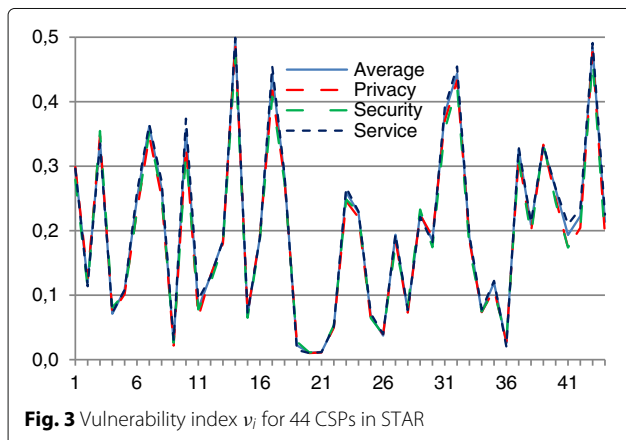
Experimental results

We prototyped CARAM and run experiments the data in STAR [4] during the course of 2015. In this section we present the results of these experiments (see Additional file 2 — CARAM Experiments.xlsx for the detailed data). Figures 3 and 4 illustrate the vulnerability index v_i and the adjusted probability β_i values for the 44 CSPs included in our analysis from STAR. These are calculated by the CARAM prototype as explained in Section “Risk levels computation”.

The differences between the vulnerability indices and therefore the adjusted probabilities of various CSPs in STAR are in the order of magnitude. For example, the lowest vulnerability index is 0.011, which is the vulnerability index of a CSP called as Iguana (we call the CSPs in STAR database by using animal names to preserve their confidentiality). On the other hand, the vulnerability index

for the highest risk CSP (i.e., Gazelle) is 0.491. Although the vulnerability index of Gazelle is more than 44 times higher than the Iguana’s, it is less than 0.5. This means that the probability value for the highest risk CSP in STAR will be reduced more than 50 %, and become “LOW” according to our higher resolution qualitative scale. This results from the fact that the CSPs from STAR declare to implement the majority of controls. That makes sense because the CSPs in STAR represent a subset of CSPs which are putting an effort to reduce their risk. Their submission of CAIQ is an indication for that.

Another interesting observation in Figs. 3 and 4 is about the differences among privacy, security and service vulnerability index and adjusted probability for each CSP. The differences among these three values are observable only for few CSPs. This is an indication that the CSPs in STAR do not focus on one of privacy, security or service, but



typically treat them equally. This is coupled with the fact that the categories of privacy, security and service risks are overlapping, and the number of risks in the privacy, security and service categories by our selection are almost the same: 19, 22 and 22 respectively (see Table 7).

For our experiments, we considered five types of CSCs based on the nature of their assets. We calculated the percentage of the risk incident scenarios that each asset is related to the ENISA risk assessment. We noticed that the following three assets are not related to any scenarios: A15 management interface APIs, A18 physical buildings, and A19 cloud provider application (source code). We call any asset related to at least 10 % of the incident scenarios as highly exposed asset. Some assets are service-related and some are data assets. Finally, data assets can be personal or not personal. Based on these, we use the following five types of CSCs in our experiments:

- CSC_1 have all the assets in the ENISA’s list.
- CSC_2 have all highly exposed assets.
- CSC_3 have all the data and service assets.
- CSC_4 have only data assets.
- CSC_5 have only personal data assets.

CARAM computes the privacy, security and service impact values adjusted for these five classes of CSCs as depicted in Fig. 5. As clearly shown, the privacy impact of risk scenarios is somewhat higher than the service and security impacts. This may be explained by the fact that most of the risks directly or indirectly may impact an asset related to privacy. Since the higher order CSCs (i.e., CSC_5 is the highest order and CSC_1 is the lowest order CSC) have only a subset of the assets owned by the lower order CSCs, their adjusted impact values are predictably lower.

In Figs. 6, 7 and 8 the risk levels for various CSC classes and three CSPs, which have the lowest, average and the highest adjusted probability, are illustrated. The darker

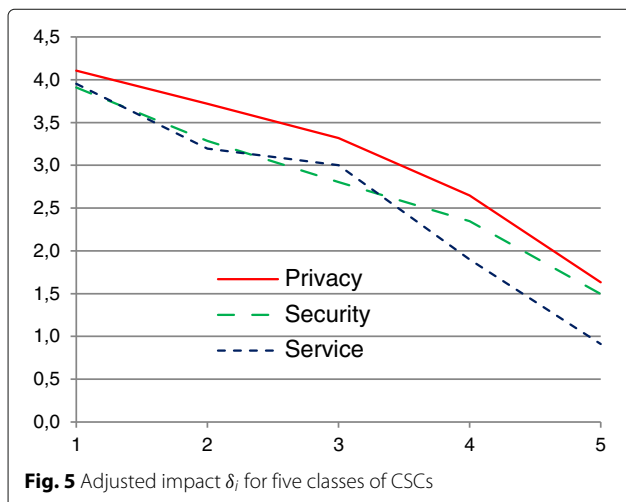


Fig. 5 Adjusted impact δ_i for five classes of CSCs

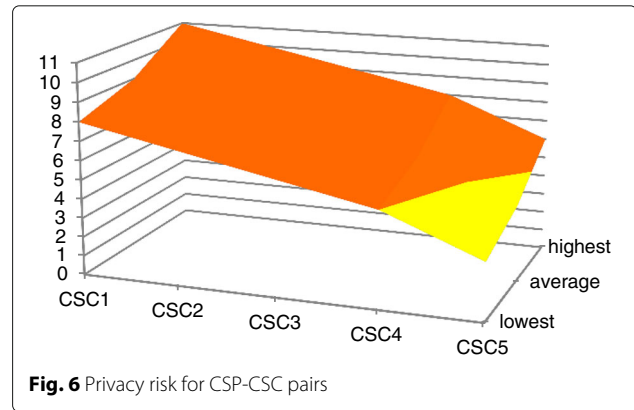


Fig. 6 Privacy risk for CSP-CSC pairs

part of the figures shows the CSC-CSP pairs with medium level of risk, and the lighter part shows the CSC-CSP pairs with low level of risk according to the ENISA’s qualitative risk scale. The vertical axis of the figures gives the risk level according to the CARAM’s higher resolution scale as in Fig. 2. We cannot observe any CSC class and CSP (the CSPs from STAR) pair that have high level of risk. This is as expected because all the CSPs in STAR apply at least half of the controls from the CSA CCM. The effects of measures are also assessed in the cloud risk assessment of CNIL [16]. Almost the same shift in risk levels is addressed in [16]. CARAM’s difference from CNIL’s risk assessment is that CARAM can take into account the effects of measures for specific CSC-CSP pairs. In Figs. 6, 7 and 8 it can be observed that privacy risk is higher than the security and the security risk is higher than the service risk. The privacy risk is low only for CSC_4 and CSC_5 when the CSPs with adjusted probability value below the average. For all the other types of CSC-CSP pairs, the privacy risk is medium.

Figures 7 and 8 indicate that the security and service risk levels for almost all the CSPs in STAR is low for CSC_3 , CSC_4 and CSC_5 . When the CSC uses cloud also for service assets (i.e., CSC_3) the security and service risks for CSPs with average or above adjusted probability becomes

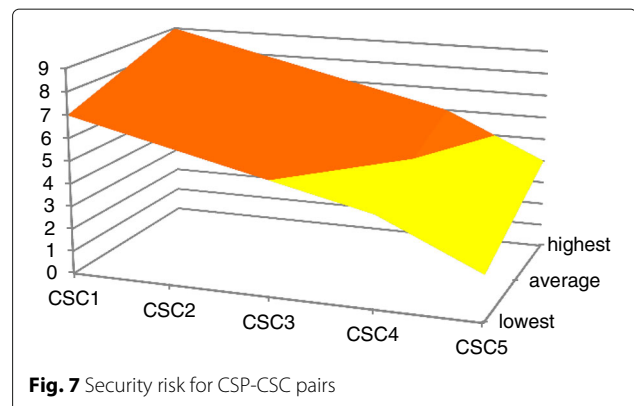



Fig. 7 Security risk for CSP-CSC pairs



medium. The security and service risks for CSC_1 and CSC_2 are always medium for the CSPs in STAR.

Limitations

There are a few limitations that may impact the accuracy of the results mainly stemming from the analyzed input data, some of which were also outlined in [5]:

- Vague formulation of the CAIQ answers provided by the analyzed CSPs: some CSPs avoid direct yes/no answers to the CAIQ questions and use generic wording instead;
- Possibility for deliberate misinformation in the CAIQ answers provided by the analyzed CSPs: CSA has a process of reporting misinformation in CSA STAR [4];
- Ineffective implementation of the security controls by the analyzed CSPs: only 3 of them have third party certification from CSA and CSA does not provide a detailed breakdown of scores for each control. To address this additional methods for evaluating control effectiveness are required, e.g. penetration testing or analysis of previous incidents (see [24, 27] for example approaches);
- In the CAIQ v1.1 [25] there is a misalignment between the description of the security controls and the actual questions that are querying their implementation; this seems to be addressed though by CSA in the newer version of CAIQ.

Notwithstanding, all results exposed in this paper are verifiable and reproducible. Variations can be introduced according to the interpretation of the cloud provider answers. We favoured an impartial, automated classification of the answers based on our training set. Smaller scale uses of CARAM can reach much better precision in the case of manual classification of the CSP answers to the CAIQ questionnaire - which is time consuming but doable and even advisable for some categories of CSCs. Such manual check of the CSPs security practices can

sometimes be done in coordination between CSP and CSC, what would lead to more transparency. As big cloud players are unlikely to invest in such reviews, CSCs can perhaps increase the pressure for (automated) continuous monitoring of security controls, which would provide more visibility to them on the security operations of the providers.

Conclusion

CARAM is a qualitative and relative risk assessment model for assisting CSCs to select a CSP that fits their risk profile the best. It is based on the existing frameworks such as ENISA, CAIQ and CNIL and complements them to provide the CSC with a practical tool. It is a risk assessment approach such that evaluation is carried out for a specific CSC, which means assessment for each CSP-CSC pair is for that pair and not generic. Moreover, the model can be easily adapted to assess further risks scenarios and/or security control frameworks, might they change in the future. These multiple features make CARAM unique with respect to the state of the art in risk assessment techniques and also among other works proposing cloud security metrics.

We have implemented a Proof-of-Concept prototype as part of the Data Protection Impact Assessment tool developed in A4Cloud project [28, 29]. The tool asks a (potential) CSC to select a CSP from a given list of around 50 providers, which answered to the CAIQ and evaluates a risk landscape of 35 risks from Table 1 grouped into 3 categories: service, security and privacy using the described methodology. Also the tool allows the CSCs to compare the risk profiles of any two providers, thus helping to select the most suitable CSP from the security point. We performed the analysis of the risk profiles for 44 CSPs from STAR and 5 imaginary classes of CSCs illustrating the coverage of security controls by the different CSPs.

Endnote

¹ The answers concerning the SAP HANA Enterprise Cloud were only available to customers on demand.

Additional files

Additional file 1: Mapping CAIQ questions to vulnerabilities. This file contains a table representing an example mapping r_{mk} of CAIQ questions to ENISA vulnerabilities (see Section The vulnerability parameter for a CSP). This is the full version of the mapping provided in Table 6. (XLSX 16.6 kb)

Additional file 2: CARAM Experiments. This file contains the simulation data (see Section Experimental results) including the risk values for several classes of CSCs for all the analyzed CSPs from STAR. (XLSX 331 kb)

Acknowledgements

This work was partly conducted during the EU-funded FP7 project titled as "Accountability for Cloud and Other Future Internet Services", A4Cloud (Grant No. 317550). We also acknowledge the FP7 EU-funded project Coco Cloud

"Confidential and compliant clouds" (Grant No. 610853) as it benefited from fruitful interactions with the project's team members.

Authors' contributions

EC drafted the model, conducted the experiments and drafted the manuscript. AG revised the model, performed a statistical analysis of the data from STAR, compiled the related work and drafted the manuscript. AS revised the model, implemented a proof of concept prototype and drafted the manuscript. YR revised the model and statistical analysis of the data from STAR and critically reviewed the manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Electrical and Computer Engineering Department, University of Stavanger, Stavanger, Norway. ²SAP Labs France, Mougins, France. ³Network Security Team, Eurecom, Biot, France.

Received: 13 August 2015 Accepted: 19 August 2016

Published online: 13 September 2016

References

- Catteddu D, Hogben G (2009) Cloud Computing. Benefits, risks and recommendations for information security. Technical report. ENISA. http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport. Accessed 13 Aug 2015
- The Notorious Nine Cloud Computing Top Threats in 2013 (2013). Technical report, Cloud Security Alliance. https://downloads.cloudsecurityalliance.org/initiatives/top_threats/The_Notorious_Nine_Cloud_Computing_Top_Threats_in_2013.pdf. Accessed 13 Aug 2015
- Article 29 Data Protection Working Party, Opinion 05/2012 on Cloud Computing (2012). http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf. Accessed 13 Aug 2015
- Cloud Security Alliance, Security CSA Trust & Assurance Registry (STAR). <https://cloudsecurityalliance.org/star/>. Accessed 25 Jul 2014
- Cayirci E, Garaga A, Santana A, Roudier Y (2014) A Cloud Adoption Risk Assessment Model. In: 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing. pp 908–913. doi:10.1109/UCC.2014.148. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7027615>. Accessed 13 Aug 2015
- ISO 31000:2009. Risk management - Principles and guidelines (2009). Technical report, ISO/IEC. http://www.iso.org/iso/catalogue_detail?csnumber=43170. Accessed 13 Aug 2015
- ISO 31010:2009. Risk management - Risk assessment techniques (2009). Technical report, ISO/IEC. http://www.iso.org/iso/catalogue_detail?csnumber=51073. Accessed 13 Aug 2015
- ISO/IEC 27005:2011 Information technology - Security techniques - Information security risk management (2011). Technical report, ISO/IEC. http://www.iso.org/iso/catalogue_detail?csnumber=56742. Accessed 13 Aug 2015
- NIST Special Publication 800-30 Revision 1: Guide for Conducting Risk Assessments (2012). Technical Report September, NIST. http://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf. Accessed 13 Aug 2015
- NIST Special Publication 800-37 Revision 1: Guide for Applying the Risk Management Framework to Federal Information Systems A Security Life Cycle Approach (2010). Technical report, NIST. doi:10.6028/NIST.SP.800-30r1. Accessed 13 Aug 2015
- COBIT 5. A Business Framework for the Governance and Management of Enterprise IT (2012). Technical report, ISACA. <http://www.isaca.org/cobit/pages/default.aspx>. Accessed 13 Aug 2015
- Stone G, Noel P (2013) Cloud Risk Decision Framework. Technical report. Microsoft. http://download.microsoft.com/documents/australia/enterprise/SMIC1545_PDF_v7_pdf.pdf. Accessed 13 Aug 2015
- Cloud Security Alliance Cloud Control Matrix (CCM). <https://cloudsecurityalliance.org/research/ccm/>. Accessed 25 Jul 2014
- Kaplan S, Garrick BJ, Kaplin S, Garrick GJ (1981) On the quantitative definition of risk. *Risk Anal.* 1(1):11–27. doi: 10.1111/j.1539-6924.1981.tb01350.x
- Ezell BC, Bennett SP, von Winterfeldt D, Sokolowski J, Collins AJ (2010) Probabilistic Risk Analysis and Terrorism Risk. *Risk Anal* 30(4):575–589. doi:10.1111/j.1539-6924.2010.01401.x
- Methodology for Privacy Risk Management; How to implement the Data Protection Act (2012). Technical report, CNIL. <http://www.cnil.fr/fileadmin/documents/en/CNIL-ManagingPrivacyRisks-Methodology.pdf>. Accessed 13 Aug 2015
- Cloud Security Alliance Consensus Assessment Initiative Questionnaire (CAIQ). <https://cloudsecurityalliance.org/research/cai/>. Accessed 25 Jul 2014
- Luna J, Ghani H, Vatea T, Suri N (2011) Quantitative Assessment of Cloud Security Agreements: A Case Study
- Luna J, Langenberg R, Suri N (2012) Benchmarking cloud security level agreements using quantitative policy trees. In: Proceedings of the 2012 ACM Workshop on Cloud Computing Security Workshop - CCSW '12. ACM Press, New York. p 103. doi:10.1145/2381913.2381932. <http://dl.acm.org/citation.cfm?doid=2381913.2381932>. Accessed 13 Aug 2015
- Habib SM, Ries S, Muhlhauser M (2011) Towards a Trust Management System for Cloud Computing. In: 2011 IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications. pp 933–939. doi:10.1109/TrustCom.2011.129. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6120922>. Accessed 13 Aug 2015
- Busby J, Langer L, Schöller M, Shirazi N, Smith P (2013) SEcure Cloud computing for CRITICAL infrastructure IT Deliverable: 3.1 Methodology for Risk Assessment and Management. Technical report, SECCRIT Project. <http://www.ait.ac.at/uploads/media/D3-1-Methodology-for-Risk-Assessment-and-Management.pdf>. Accessed 13 Aug 2015
- SPECS Secure Provisioning of Cloud Services Based on SLA Management. <http://www.specs-project.eu/>. Accessed 13 Aug 2015
- Luna J, Taha A, Trapero R, Suri N (2015) Quantitative Reasoning About Cloud Security Using Service Level Agreements. *IEEE Transactions on Cloud Computing* PP(99):1–1. doi:10.1109/TCC.2015.2469659
- Cayirci E (2013) A joint trust and risk model for MSaaS mashups. In: 2013 Winter Simulations Conference (WSC). pp 1347–1358. doi: 10.1109/WSC.2013.6721521. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6721521>. Accessed 13 Aug 2015
- Cloud Security Alliance Consensus Assessments Initiative Questionnaire V1.1. <https://cloudsecurityalliance.org/download/consensus-assessments-initiative-questionnaire-v1-1/> Accessed 25 Jul 2014
- Machine Learning Group at the University of Waikato WEKA 3: Data Mining Software in Java. <http://www.cs.waikato.ac.nz/ml/weka/>. Accessed 25 Jul 2014
- Habib SM, Varadharajan V, Muhlhauser M (2013) A Trust-Aware Framework for Evaluating Security Controls of Service Providers in Cloud Marketplaces. In: 2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications. pp 459–468. doi: 10.1109/TrustCom.2013.58. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6680875>. Accessed 13 Aug 2015
- Cloud Accountability Project. <http://www.a4cloud.eu/>. Accessed 10 Aug 2015
- Garaga A, Santana de Oliveira A, Cayirci E, Dalla Corte L, Leenes R, Mhungu R, Stefanatou D, Tetrimida K, Felici M, Alnemr R, Pearson S, Vranaki A (2014) D:C-6.2 Prototype for the data protection impact assessment tool. Technical report, A4Cloud Project. <http://www.a4cloud.eu/sites/default/files/D36.2%20Prototype%20for%20the%20data%20protection%20impact%20assessment%20tool.pdf>. Accessed 13 Aug 2015