Journal of Universal Computer Science, vol. 18, no. 12 (2012), 1662-1678
submitted: 15/9/11, accepted: 15/2/12, appeared: 28/6/12 © J.UCS

Analysing the Security Risks of Cloud Adoption Using the SeCA Model: A Case Study

Thijs Baars
(Department of Information and Computing Sciences, Utrecht University, The Netherlands
t.baars@uu.nl)

Marco Spruit
(Department of Information and Computing Sciences, Utrecht University, The Netherlands
m.r.spruit@uu.nl)

Abstract: When IS/IT needs to be replaced, cloud systems might provide a feasible solution. However, the adoption process thus far has gone undocumented and enterprise architects are troubled with proper hands-on tools missing, until very recently. This single case study describes a large Dutch utility provider in their effort to understand the facets of the cloud and identifying the risks associated with it. In an action research setting, the SeCA model was used to analyse the cloud solutions and identify the risks with specific data classifications in mind. The results show how decision makers can use the SeCA model in various ways to identify the security risks associated with each cloud solution analysed. The analysis assumes that data classifications are in place. This research concludes that by using the SeCA model, a full understanding of the security risks can be gained on an objective and structural level; this is a further validation of prior empirical research that the SeCA model is a proper hands-on tool for cloud security analysis.

Keywords: SeCA Model, Cloud Security, Cloud Computing, Case Study, Information Security
Category: H.2, H.3.7, H.5.4

1 Introduction

When deciding to adopt cloud architectures within an organization, security issues are cited as the primary issues for decision makers that withholds a positive outcome [Foster, Zhao, Raicu and Lu 2008; Ghinste 2010; Mowbray and Pearson 2009]. There is good reason for this, security issues in the cloud can be very complex; looking at the definition of the cloud, as defined by NIST and ENISA, the cloud consists of three deployment models, four delivery models and eight characteristics [Hogben and Catteddu 2009; Mell and Grance 2010]. All these models and characteristics can influence the security risks and threats of the cloud [Baars & Spruit 2012].

Furthermore, enterprise architects are troubled with proper hands-on tools missing, until very recently. Even though a variety of tools is proposed [Ko, Jeon & Morales 2011; Li, Yang, Kandula & Zhang 2010; Popa, Yu, Ko, Ratnasamy & Stoica 2010] none of them have been documented as being used in practice.

In this research we took one of those tools, the SeCA model [Baars and Spruit in-press, 2012], into practice to understand how well it performs in the hands of an implementer. The goal is to verify the assumptions made in developing this tool, as well to affirm its practical usability in the work field.

Does this model help decision makers analysing cloud services as it is designed to?

The SeCA model [Baars and Spruit in-press, 2012] is a model that allows decision makers and IT professionals to analyse cloud services and architectures on their security specifications by using data classifications as a reference. These data classifications are assumed to be already in place for data storage and computations systems currently in use by the organization. This model has been verified by an expert panel during a delphi method [Baars and Spruit in-press], but has never been put into practice. This research therefore validates the SeCA model by using it in practice.

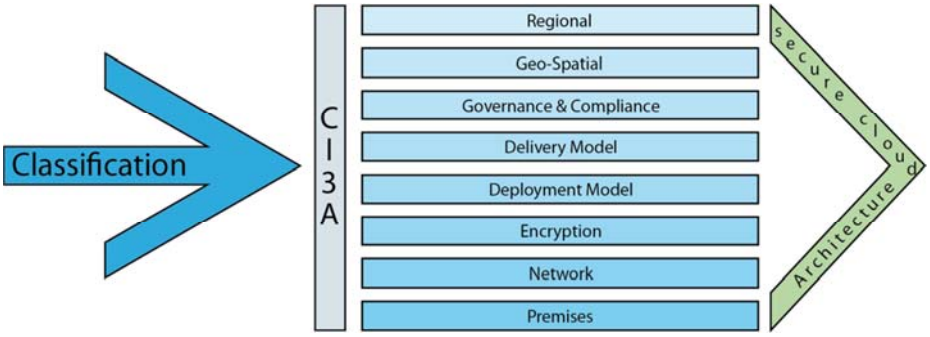

Figure 1: The SeCA Model [Baars and Spruit in-press]

The SeCA model assumes that data is classified in several classifications (usually four). The data storage and computing systems in use have different security specifications to which they have to adhere to count as a secure solutions for a specific data classification. This allows for secure processing and storage of data in each data classification. In practise this may mean that multiple data storage systems are in place, specific to a certain data classification in order to adhere to the stricter security guidelines that come with a stricter (or higher) data classification.

The difference in data classifications is depicted in the model as the arrow on the far left. From there, every solution is viewed within the scope of the CI3A, Confidentiality, Integrity, Availability, Accountability and Auditability.

The CI3A is an extension of the *de facto* CIA triad. The SeCA model utilizes CI3A to maintain the right level of assurance within the environment. The CIA-triad does not cover the complexities of the cloud, as experts from the Baars & Spruit delphi study [in-press] mentioned and recent research has shown [Hu and Xu 2009; Mowbray and Pearson 2009; Peterson and Gondree 2011; Wang and Zhou 2010]. The many variations of the cloud, a total of 72 (3 *delivery models* × 3 *deployment models* × 8 *Characteristics*), as defined by NIST [Mell and Grance 2010] and excluding any newly developed deployment models, such as xCloud [Williams et al. 2011] created a complexity which led to the CI3A [Baars and Spruit in-press].

They CI3A is a frame of reference when analysing the attributes (shown as horizontal bars in the mode.) For example when analysing the premises attribute of a

cloud architecture, one investigates how auditability is affected by the fact that the researched cloud architecture is on or off premise.

Essential in the SeCA model is the realization that physical location is of mere importance. The physical location influences at least four of the eight attributes defined in the model. These include Regional (in which jurisdiction are the servers located), geo-spatial (how far apart are the servers from each other, what physical aspects do the datacentres have), Governance & Compliance (to which rules and regulations does the cloud solution adhere to?) and premises (are the servers located on organizational premises or not?)

This paper describes the research of a single case study conducted at a large Dutch utilities firm. In the next section, the related literature that influences this research is described. In the "Research Methods: The Case Study At The Utilities Firm" section, the research method, the case study outline and the utilities firm are described. In "The SeCA Model" the SeCA model is explained in more detail, which is the tool of use in this research. Following that is the results section, which is split in half. "Applying the SeCA Model at the Utilities Firm" discusses the first results, elaborates on the process and data aggregation. It clearly shows how the SeCA model was used and how practice meets theory in this case study. "Mapping Data Classifications & Ranking the Stars" shows the next step in data analysis. It describes how the data classifications are mapped to the candidate cloud solutions. It results in a rank of best-fit cloud solutions to the utilities firm. Next, we discuss the conclusions from this research, with possibilities for further research.

2 Related Literature

In 2011, the SeCA Model was developed by Baars & Spruit [in-press, 2012], in an effort to provide decision makers and IT professionals with a tool to analyse cloud architectures. It is based on existing literature in three ways, first some models were already in place or developed when the SeCA model was developed, these are outlined below. Second, Tools for cloud computing have influenced the SeCA model, outlined after the models subsection (in 2.2). Third and last, innovations in cloud computing have influenced the discussion towards the model as well as the model itself. These are outlined at the end of this section (2.3).

2.1 Models

The SeCA model was meant as an upgrade from the Jericho Forum's cloud cube model [Jericho Forum 2009]. The Cloud Cube Model was found to be an oversimplification of the risks and threats within the cloud and thus not capable of giving users a comprehensive overview of cloud architectures. It consists of four attributes instead of the eight in the SeCA model.

[Siebenhaar, Tsai, Lampe & Steinmetz 2011] describe a holistic model for analyzing and modeling security aspects of cloud-based systems. The granularity of the SeCA is finer, although most attributes of the SeCA model are implied by this model.

[Almorsy, Grundy & Ibrahim 2011] provide a collaboration based framework for determining security management. The presented SeCA model in this research differs

from the Almorsy et al. model as it focuses on security measures within the architecture of the cloud determined by data classifications.

Next to the cloud cube model and the SeCA model, [Kaliski Jr & Pauley 2010] offer their thoughts on the paradigm of risk assessment as a service; an automated way of risk assessment of cloud services.

[Takabi, Joshi & Ahn 2010] provide an overview of challenges within the security and privacy of cloud environments. The SeCA model discusses them as well, and models them in a user friendly framework, where Takabi et al. just sums them up.

The SeCA model adds granularity as well as a hand-on tool for decision makers, implementers and IT professionals for assessing cloud security.

2.2 Tools

CloudCMP developed by [Li et al. 2010] is an effort to provide end-users with a framework for selection of a cloud provider using benchmark results and select criteria. It could be used as a marketplace for generic SeCA model outcomes, showing which solution offers which security options as part of their service. It is in line with [Krautheim 2009] which gives users also more control, but on an architectural level where it user editable security controls of cloud solutions in virtual datacentres.
[Popa et al. 2010] introduce CloudPolice, a system at takes security down to the hypervisor level.

[Itani, Kayssi & Chehab 2009] developed what they call “Privacy as a Service”, a service that uses cryptographic co-processes to ensure secure data processing within a cloud environment. [Troncoso-Pastoriza and Pérez-González 2010] propose CyptoDSP, a secure way of signal processing in the cloud, [Wood et al. 2010] propose to use cloud service as disaster recovery, creating disaster recovery as a service.

[Kaliski Jr and Pauley 2010] argue that risk assessment as a service could be viable to understand risks and threats within cloud environments. They argue that SLAs are key. [Benson, Akella and Maltz 2009] also investigate SLA policies.

The research described above can directly applied to measures within the SeCA model, or has influenced it. It is clear that users should get more control in cloud security and these tools may provide developments or implementations for such controls.

2.3 Innovations

[Hu and Xu 2009; Peterson and Gondree 2011; Tiwana, et al. 2010] provide proof that location of cloud providers is affecting performance and security of cloud solutions. This aspect isn’t mentioned in the cloud cube model but is in the SeCA model.

[Wang, Wang and Ren 2009] are using homomorphic keys to ensure data reliability, and state that cloud security is still in its infancy [p.9], a vision the authors of this paper can only adhere to. [Ibrahim, Hamlyn-harris and Grundy 2010] provide an overview of emerging security challenges and propose research challenges that come along with it. They propose issues with mapping security issues to virtualized environment, something the SeCA tries to accomplish.

[Pal, et al. 2011] provide a new framework for cloud users to identify collaborative and trusted users. It works on the infrastructure level and helps to prevent unauthorized access to user's data. [Wang & Zhou 2010] developed a mechanism to create multitenant clouds accountable.

All these innovations showed that there is a need for a secure cloud and that once the issues a clear, by using for example the SeCA model, solutions are plenty to secure your cloud.

3 Research Methods: The Case Study at The Utilities Firm

This single case study [Benbasat, Goldstein and Mead 1987; Yin 2009] describes a large Dutch utility provider in their effort to understand the facets of the cloud and identifying the risks associated with it.

In an action research setting the SeCA model is used to analyse the cloud solutions and identify the risks with specific data classifications in mind. Action research was particularly useful for this case, as the researchers could participate in solving the issues of the utilities firm, while maintaining the original intent to validate the model [Checkland 1981; Susman and Evered 1978].

The research was setup as follows: First a meeting was held to get an overview of the firm, the IS/IT strategy, the goals and expectations of cloud adoption. Next, four sessions lasting a day each were held in which

1. *A list of cloud services was gathered from various sources.* These sources included Gartner reports, industry magazines and search results from the Google search engine. The services found were listed in a template file (henceforth called the matrix) for the SeCA model. This template file, included in [Baars and Spruit in-press] has fields for all the attributes of the SeCA model. In this step however, the service name and service provider were only listed. An example of this matrix is shown in table 1.
2. *Cloud services were analysed using the matrix that was derived from the SeCA model.* Information was gathered using a variety of sources, the original service provider website and salespersons being the most important. Information found was added to the matrix where appropriate.
3. *Data classifications are mapped to the SeCA model.* These were adapted for the cloud and received from the Chief Information Security Officer on forehand. In essence this is a new row in the matrix per data classification, which shows which security measures are needed to adhere to that data classification. This way, one can easily compare the analysed cloud service and the minimum specifications required for a certain data classification.
4. *A ranking algorithm is empirically formulated based on the data classifications.* This ranking algorithm ranks the services so that an objective score list can be created for the to-be selected cloud service. This is more detailed in section 5.2.
5. *The cloud services and data classifications are mapped to each other to short-list and rank the services.* In an effort to understand the results of the analysis and to see clearly which cloud services has the best fit with which data classification, a short list is created from the in step 4 ranked services. This process yields a short list of services that can be further examined (for

feasibility analysis and other analyses not part of the SeCA model) or presented as the final outcome. This step is detailed in section 5.2.

As a final step, the results are presented to the management team responsible for the ultimate decision in cloud adoption. Unless noted otherwise, all meetings were held with one and the same enterprise architect of the utilities firm.

In all meetings, the researchers discussed the steps undertaken in this process, and helped at certain points to speed-up the process. For example, when retrieving information to fill in the matrix on for example the geographic locations of datacentres, the researchers would try to find that information while the enterprise architect would direct its attention on gathering encryption methods in place. As a part of validation that information gathering can be done by an IT professional as the enterprise architect in question, a service was totally analysed by both the researchers and the enterprise architect. The only difference found was "the stunning amount of acronyms" the author used in comparison of the analysis results by the enterprise architect. It was at that moment decided that combined information gathering would not influence the results and would in fact recreate a group project setting common at the organization.

3.1 The Utilities Firm

The firm at which the case study was conducted is one of the largest utilities companies of the Netherlands, with over 2.6 million connected households, organizations and governments and net results totalling 193 million euros. Being PAS 55-1 and ISO 9001 certified, security is in the core of the more than a century old company. As an outage in service equals to an outage in the power supply to the households and companies, business decisions are made carefully. In 2010 a major IS/IT transition was made, forcefully, as the Dutch government liberated the energy market. This meant that many utility companies in The Netherlands, including the one discussed here, were split in two separate organizations.

For this reason, IS/IT systems are swiftly being disconnected from the former mother-firm in an effort to become a standalone organization.

The decision has been made to investigate opportunities in the cloud.

The project leader for this investigation is a senior enterprise architect at corporate level. In an effort to grasp the complexities of the cloud and the services it could provide for the organization, he turned to Utrecht University as he was made aware of the development of the SeCA model. It was decided that for three different product categories an analysis would be performed: "office suites", "database platforms" (or database as a service) and "identity and access management services". For the briefness of this paper, we will only discuss the office suites. The process for the other service categories is similar to the one described in this paper.

4 The SeCA Model

The SeCA model (figure 1) gives an abstract overview of all the characteristics of the cloud as defined by NIST [Mell and Grance 2010] and ENISA [Hogben and Catteddu 2009] extended with the CI3A and attributes derived from the delphi study help for the development of the SeCA model [Baars & Spruit in-press].

The model assumes a data-centric approach, it does not discriminate between the types of information stored nor the file formats; it does discriminate in the user rights of the data stored and processed using its attributes explained below. The model therefor uses data classifications as data classification describe who can and cannot see, use and execute data, and under which circumstances. These data classifications should be defined and implemented on forehand.

When using the model, a data classification is inputted and analyzed, or dissected if you will, using the attributes in the model which are displayed as horizontal bars. The model outputs guidelines for the cloud environment and to which specification a cloud solution should adhere on the basis of data classifications. Concrete, this will probably be a list of specifications which comprise of a minimum requirements for which cloud services should adhere to be able to accept, store and process data that classified under the just analyzed data classification.

However, the model can be used in two directions, from left to right (forward direction), or from right to left (backward direction). The following section describes how one can use this model in both directions.

4.1 The Forward Direction

The forward direction, seen in figure 1, from left to right, takes a data classification, as discussed above, an input, and assesses all attributes on the basis of that data classification. Thus creating an extension for the data classification assessed. For each classification in place in the organization, this routine will be executed. The goal of this method is to create a set of classifications that is ready for the cloud. Thus, creating a list of extremes to which a cloud architecture has to adhere in order to be applicable to the data classification.

4.2 The backward Direction

The backward direction was not an original intent, but was implicated by the Baars & Spruit [2012] paper. It is the method of use in this case study, by choice of the enterprise architect. Instead of taking a classification as the input, it takes cloud solution as the input (in image 2, the right side), and then retrieves the attributes that are specified in the model. This retrieval can be as simple as obtaining the information from a technical specification sheet. These are then compared with the classification in place (which might need to be adapted for the cloud, as many classifications do not specify rules comprising the complexities of the cloud; left side)

From all the solutions analyzed a shortlist can be created of services that are viable for adoption, shown in the arrow pointing downwards.

In this case study, these solution were presented to the management team. As four data classifications were in place in this case study, the backward direction was proven to be fruitful. By analyzing the possible solutions first, a clear overview was created of the environment in which all solutions operated. By then comparing these environments with the architectures defined by the classifications, one can easily distinguish the solutions that are viable for adoptions within the norms set forth by the classifications. This process is described in detail in the next section.

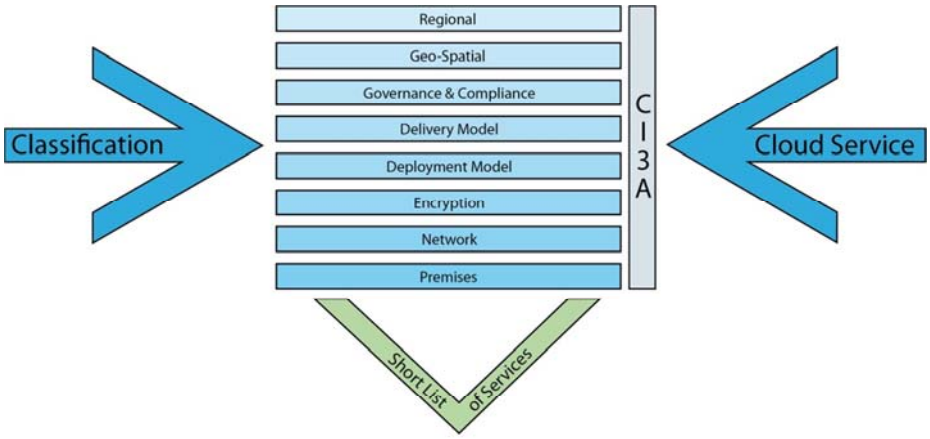

Figure 2: Backwards direction of the SeCA model. Analyzing cloud services for comparison with data classifications, resulting in a short list of Cloud Services.

5 Results

This section describes the findings from the case study. It is split into two section, an "Applying the SeCA Model at the Utilities Firm" and a "Mapping Data Classifications and Ranking the Stars" section. The former introduces the aggregation of data for the shortlisting of the cloud solutions. These solutions are then analysed. This analysis follows the attributes from the SeCA model. The latter section maps the existing data classification that are already in place at the utilities firm, to the SeCA model attributes. From the se classification a ranking algorithm is then computed, which in turn creates a ranking of all cloud solutions and how well they fit to the security classifications within the firm. This allows for a structural, objective analysis.

5.1 Applying the SeCA Model at the Utilities Firm

When we started to use the SeCA model, the enterprise architect chose to use the backward direction as it fitted with the question the firm was struggling with: "What applications could be outsourced to the cloud?" As application, or services, are centric in this question, using the backward direction seems trivial, but it also allowed for the Chief Information Security Officer to adapt the data classifications in place to the complexities of the cloud (see below for a more detailed description.)
At the start of the analysis, three major office suite offerings were found: SuiteA, SuiteB and SuiteC.

To perform the analysis using the SeCA model, retrieving the specifications of the cloud architecture, a matrix was created. This matrix combines all the attributes of the SeCA model, extended with properties of these attributes as described by earlier research [Baars & Spruit, 2012].

Each service is displayed on a separate row; the values for all the columns are then subsequently filled. The information for these values can come from white papers, marketing publications and other data sources describing the service. In this case, data sources were limited to publication by the service provider in order to prevent any misinformation from third parties. This can be tedious work, as many companies do not provide the information needed with ease. Sometimes no information could be found, nor where salespersons willing to provide that information. We will leave it in the middle what that says about the company and its security policy. See Table 1 for the matrix.

Note that the information in the table is truncated for overview and ease of understanding the process. Yet, even without truncation, the information retrieved is sparse and vague in certain points. SuiteA bluntly states that its servers are stored worldwide, without any further specification, SuiteC states that for European customers, it's data is hosted within the jurisdiction of the EU, but that it does mirror the data to US servers. Legally, this means that the data is stored in both EU and US territories. Particularly vague are the audits that take place. SuiteC simple stated that it does audits. No more, no less. SuiteB was not able, nor willing, to provide any information at all, SuiteA states it undertakes global audits, whatever that may mean.
Even though these results seem at some points grim, it must be noted that in further investigations following the analysis of the SeCA model more information may be retrieved. This can happen through personal contact with sales team during sales negotiations.

5.2 Mapping Data Classifications & Ranking the Stars

The data classifications, which were already in place, were adjusted for the cloud. The Chief Information Security Officer (CISO) used the SeCA model as a template to provide the allowed ranges of values for each of the attributes in the SeCA model and thus the cloud architecture. Table 2 shows the mapped classifications, an 'x' shows a permissible option.

These classifications were then mapped to the analysed cloud solutions.

Interesting to note was that on first sight, it was clear that SuiteC was possibly not suited for adoption as jurisdiction and location are very ill defined. Whereas the other suites could provide location up to the city where the datacentre is located, SuiteC was not willing to provide this information. However, because the manufacturer of SuiteC provided already for many software packages within the firm, there was a strong biased view in favour of SuiteC.

In an effort to overcome this bias, the classifications were mapped to the cloud solutions. This mapping took place in the form of a ranking. By creating a formula that defines the coherence of the classifications on the cloud solutions, an index number could be generated that would then be ranked (higher being better).

These impacts, or weights, were defined on an empirical basis on company culture. Being not very scientific as it is, it is as the enterprise architect mentioned, "the way most decisions are being made within the firm".

Service provider	**Deployment model**	**Delivery Model**	location			Security		
			jurisdiction	**premise**	**Network**	**Backups**	**location**	**A&A****
SuiteA	Public	SaaS	safe harbour, Worldwide	out	Both	synchronization, mirroring & fragmentation and DR*	worldwide	LDAP, user/pass, custom 2factor
SuiteB	Public	SaaS	Safe harbour , USA, NJ & CA	out	out	mirroring, DR*	USA, worldwide	SAML (LDAP/AD etc.)
SuiteC	Public	SaaS	EU & USA mirroring	out	out	synchronization in USA, backup in EU	Amsterdam & Dublin	AD setup, user/pass login

*Table1: The Matrix used for the SeCA Analysis. . *DR = Disaster Recovery, ** A&A = authentication & authorization*

Service provider	Security (continued)		CI3A				
	Encryption	**certification**	**Confidentiality**	**Integrity**	**Accountability**	**Auditability**	**availability**
SuiteA	SSL, TLS	SAS 70-II, SOX, PCI-DSS	no dedicated hardware, proper HRM strategy	SSL, market-place apps	full logging, data belongs to end-user, 3rd party's	global internal audits, unclear	99.9% uptime,
SuiteB	256bit SSL, TLS	TrustE	no dedicated hardware, no info on HRM strategy	SSL	end-user = responsible	Undefined	99.9% excl. scheduled downtime
SuiteC	128bit SSL/TLS	SAS 70-II, ISO 27001	dedicated hardware possible	SSL	end-user = responsible	yes, but unclear	99.9% uptime

Table 1: continued.

Classification	Deployment model			Delivery Model			Location			security	
	public	**Comm**	**Priv**	**IaaS**	**SaaS**	**PaaS**	**jurisdiction**	**premise**	**Network**	**Backup**	**location**
Public	x	x	x	x	x	x	worldwide	in & out	in & out	Yes, mirroring/synchronisation satisfies	worldwide
Private	x	x	x	x	x	x	EU	in & out	in & out	Yes, mirroring/synchronisation satisfies	worldwide
Secret	x	x	x	x	x	x	EU	in & out	in & out	Yes, mirroring/synchronisation satisfies	EU
Top Secret		x	x	x	x	x	EU	In	In	Yes, mirroring/synchronisation satisfies	EU

Table 2: data classification mapped to the SeCA model, the 'x' signifies an allowed state

Classification	**Security (continued)**			**CI3A**				
	A&A	**Encryption**	**certification**	**Confidentiality**	**Integrity**	**Acountability**	**auditability**	**availability**
Public	user/pass for integrity	No	no	no	Yes (SSL)	no	no	Yes, brand management
Private	Employee only access	No	depending on location	Yes, employee only	Yes (SSL)	logging	depending on location	Yes, processes & employee access
Secret	Yes	Yes	depending on location	Yes, elevated employees + partners	Yes (SSL)	full logging	depending on location	Yes, mission critical
Top Secret	2factor	Yes	depending on location	Yes, special elevation employees	Yes (SSL)	full logging	depending on location	Yes, strategically critical

Table 2: continued, Encryption depicts special needs for encryption; authentication processes and such should be encrypted, but this is an industry standard.

First of all, weights were assigned to the categories in the matrix (security, delivery model, deployment model, location, CI3A) from 0.0 to 1.0. Then, from the data classification possible values for the attributes (IaaS, PaaS, backups, etc.) are assigned, likewise on a 0.0 to 1.0 scale. These values were given in respect to the classification.

Null values were assigned if there is no impact on the classification, so that they would not influence the index calculation.

In other cases kill values could be defined. If a certain value is equal to the kill value, the solution would not be viable for adoption.

If for example, a deployment model was used which isn't allowed according to the classification, this would be a kill factor and the solution would be scrapped from the list, or ranked last. If a SaaS delivery model was allowed and the service is a SaaS application, a full point can be given, stating that it meets the criteria off the classification.

Weights between the extremes were given on a relative basis where open values were possible, such as the authentication protocol; more freedom and stronger encryption than prescribed by the classification would be given a higher point than solutions meeting the minimal criteria of the classification.

This process is depicted in the following formula:

$$\begin{aligned} RankIndex = {} & weight_{deliveryModel} * (value) + weight_{deploymentModel} * (value) + weight_{location} \\ & * (value_{juridiction} + value_{premise} + value_{network}) + weight_{security} * (value_{backup} \\ & + value_{location} + value_{A\&A} + value_{encryption}) + weight_{CI3A} * (value_{confidentiality} \\ & + value_{integrity} + value_{availability} + value_{accountability} + value_{auditability}) \end{aligned}$$

After the appropriate weights are defined, and each solution has been given its values, the *RankIndex* numbers are calculated and ranked, where the highest *RankIndex* score is the best fit with the data classifications in place and thus scores the highest.

As more than one data classification is in place, multiple ranks will be made; for each classification one. This provides a benefit in trying to create an overview which solutions are applicable for data with a certain data classification (and thus which solutions can be used with all data classifications.) Each classification might have different values for each attribute.

Displayed below are the chosen weight for the categories (which are the same for each classification) and the values of the attributes within each category. These can be different per classification, and are therefore displayed as a row of 4 numbers, comma-separated. The first number represents the value for the public classification, the second the value for the private classification, the third for the secret classification and the last number represents the value for the top secret classification.

In case there is a 0 for weight, it means that the value has no impact on the choice. A kill factor is applicable when the value does not correspond with the classification's allowed values.

For the utilities firm, depicted in table 3a and 3b are the weights and values that were assigned to SuiteC.

category	Weight	Attribute	Value	Description
Delivery model	0.2	Type	0,0,0,0	Delivery model has little impact on its own, but if it does not meet the criteria of the classification it is a kill factor
Deployment model	0.2	Type	0,0,.5,.5	Deployment model has little impact on its own, but if it does not meet the criteria of the classification it is a kill factor

Table 3a: delivery model and deployment model weights explained

Category	Weight	Attribute	Value	Attribute	Value	Attribute	Value
Location	0.6	Juris-diction	*0,0,0,0*	Premise	*0,0,.5,.5*	Network	*0,0,.5,.5*
As many values had no direct impact on the classification, a lower weight was chosen than for example security. Notice that network and premise is dependent on delivery and the deployment model which according to the classification only have kill factors and are thus indirectly indifferent. If jurisdiction is outside the EU, it is a kill factor for all but public.							

Security	1.0	Backup	*.2,.5,.9,.9*	Location	*0,0,.4,.4*	A&A	*1,1,1,1*
		Certification	*0,.8,.8,.8*	Encryption	*0,0,.8,1*		
As security is the core of the firm, it was given the highest weight. Being a semi-public firm, location is of less importance and to whom the data is shared, most of it is public anyhow is the perception. However, who can change the data was found to be of utter importance and thus Security was given a high weight. Location has a kill factor if it is outside the EU in the secret and top secret classifications							

CI3A	0.8	Confidentiality	*0,1,1,1*		Availability	*.9,.9,1,.9*	
		Accountability	*0,.5,.8,1*	Auditability	*0,.2,.6,1*	Integrity	*.8,.9,.9,.9*
As the CI3A is in direct correspondence with security, it was given a weight of 0.8; CI3A also defines certain functions of data in the firm, for which the impact can be huge if that data does not meet the CI3A criteria.							

Table 3b: location, security and CI3A weights explained.

If we compute the rankings from the given values with their weights in each categories the following ranking is as displayed in table 4 below. For conciseness, the values of the SuiteA and SuiteB are not fully displayed, but only the calculated rankIndex scores are displayed.

Ranking public	Solution Name	RankIndex
1	SuiteA	3.75
2	SuiteB	3.10
3	SuiteC	2.56

Table 4a: Ranking for the pubic classification

Ranking private	Solution Name	RankIndex
1	SuiteA	7.64
2	SuiteC	5.19
3	SuiteB	4.48

Table 4b: Ranking for the private classification

Ranking secret	Solution Name	RankIndex
1	SuiteA	7.67
2	SuiteB	6.88
3	SuiteC	KILL

Table 4c: Ranking for the secret classification

Ranking top secret	Solution Name	RankIndex
1	SuiteC	KILL
2	SuiteA	KILL
3	SuiteB	KILL

Table 4d: Ranking for the top secret classification

As can be seen in table 4d, none of the suites meets the criteria for the Top Secret classification. This classification does not allow data to be stored or processed in a public cloud architecture, thus resulting in a kill factor for all Suites as they all offer their services in a public cloud. As all office suites are hosted on a public cloud, computing a ranking is barred by this kill factor. This is also shown for suiteC in the secret classification, due to the lack of a rankIndex, it is placed at the bottom of the ranking. Here, the data is mirrored to the United States without a safe harbor certification. According to the classification (and EU law) this is unacceptable.

6 Conclusions & Further Research

This research shows how the SeCA model is used in practice. Not only does it show that the SeCA can be a valuable tool in the decision making process whether or not certain data can be outsourced into the cloud, with the added ranking algorithm it is possible to get an objective overview of services provided. We therefore conclude that by these means of empirical validation, the SeCA model is a proper, useful and correct tool for the analysis of cloud architectures.

As mentioned above, there was a positive bias for SuiteC as more software packages of SuiteC were already in place in the organization. However, as shown in Table 4, SuiteC does not compete as well as the other suites. Not only doesn't it meet the specifications for storage of data classified as Secret or Top Secret, it tends to underperform in the index overall.

When we presented these results to the management team, they were grateful that such an elaborate overview was presented yet still easy to discern which solutions were capable for which classifications. The ranking algorithm enables decision makers to quickly evaluate the analysis and make the correct decisions for the organizations, whereas the professional can still see the in-depth process and understand the intricacies of the analysis.

The results show how decision makers can use the SeCA model in various ways to identify the security risks associated per cloud solution per data classification. This research concludes that using the SeCA model, a full understanding of the security

risks can be gained objectively and on structural level; a (further) validation of prior empirical research that the SeCA model is a proper hands-on tool for cloud security analysis. As the model provides an objective view of the security threats that may occur, subjectivity, such as measured with familiarity for certain brands, can be overcome. The ranking formula provides a clear overview of the best fitting services to the data classifications in place. It can clearly shortlist the investigated services for a deeper inspection on other grounds such as strategic fit and user acceptance. By using a backward direction, instead of the forward, new software suites can now easily be analysed.

This analysis obviously takes the security aspect only into account. A next step is to see how the short listed systems fit within the current environment; feasibility studies of specifics like usability, financial cost and strategic fit are also in need to be executed. However, apart from a security analysis tool, in this case study the SeCA model has also proven itself as a tool for shortlisting.

As mentioned, getting the information for the matrix can be hard. Often company representatives are not willing to disclose exact locations, and other details of the services provided. Remarkably, these locations can be found in public documents such as whitepapers. Often news articles were of much use, but these need to be thoroughly checked. Blog posts by experts and other unofficial sources of information often have discrepancies in their posts. However, confronting staff from the cloud provider with that information has proven to be successful technique in getting the right information. It is peculiar that service providers are so sparsely with their information, in the end that information could generate new clients.

The execution and development of the model has now all been done by manual labour, during the coordinated sessions and during hours aside from the meetings. This has proven to be a time exhaustive process that could be automated. We therefore would like to propose a system that persistently stores the specification of different services through the means of vendor input or crowd sourcing. A user could then select the category of services it is willing to analyse and get a quick and clear overview of services that may be applicable. The user could provide its own weight/formula and create a ranking on the fly. Furthermore, the most time consuming process when doing it by hand is the aggregation of information. This information will be the same for every organization. The cloud solutions in the end do not change per user. By automating this, the SeCA model can become a very quick tool to analyse cloud solutions whilst keeping the precision needed for IT pro's and security experts.

References

[Almorsy, Grundy and Ibrahim 2011] Almorsy, Grundy, J. & Ibrahim, A. S. "Collaboration-Based Cloud Computing Security Management Framework"; Proceedings of the 2011 IEEE International Conference on Cloud Computing (CLOUD2011), IEEE, Washington, DC (2011) 364–371

[Baars and Spruit in-press] Baars, T. & Spruit, M. R. "The SeCA Model: Ins & Outs of a Secure Cloud Architecture". In D. Rosado, D. Mellado, E. Fernandez-Medina, & M. Piattini (Eds.), Security Engineering for Cloud Computing: Approaches and Tools. IGI Global, Hershey

[Baars and Spruit 2012] Baars, T. & Spruit, M. R.. "Designing a Secure Cloud Architecture: The SeCA Model"; International Journal of Information Security and Privacy, 6, 1 (2012) 14-32

[Benbassat, Goldstein and Mead 87] Benbasat, I., Goldstein, D. K. & Mead, M. "The Case Research Strategy in Studies of Information Systems"; MIS quarterly, 11, 3 (1987) 369-386

[Benson, Akella and Maltz 2009] Benson, T., Akella, A. & Maltz, D. A. "Mining policies from enterprise network configuration"; Proceedings of the 9th ACM SIGCOMM conference on Internet measurement ACM, Chicago, IL (2009) 136-142

[Checkland 81] Checkland, P. "Systems Thinking, Systems Practice", John Wiley and Sons, New York (1981)

[Foster et al. 2008] Foster, I., Zhao, Y., Raicu, I. & Lu, S. "Cloud computing and grid computing 360-degree compared" Proceedings from the Grid Computing Environments Workshop, IEEE, Austin, TX (2008) 1-10

[Ghinste 2010] Ghinste, B. "Gartner: Private Cloud Computing Plans From Conference Polls"; MSDN Blogs (2011) Retrieved from http://blogs.msdn.com/b/architectsrule/archive/2010/05/07/gartner-private-cloud-computing-plans-from-conference-polls.aspx

[Hogben and Catteddu 2009] Hogben, G. and Catteddu, D. "Cloud Computing: Benefits, Risks and Recommendations for Information Security"; ENISA Report 51273, Crete, Greece (2009)

[Hu and Xu 2009] Hu, H. and Xu, J. "Non-Exposure Location Anonymity"; Proceedings of the 25th International Conference on Data Engineering, IEEE, Shanghai, CN (2009) 1120-1131

[Ibrahim, Hamlyn-harris and Grundy 2010] Ibrahim, A. S., Hamlyn-harris, J. H. & Grundy, J. "Emerging security challenges of cloud virtual infrastructure"; Proceedings of the APSEC 2010 Cloud Workshop, IEEE, Sidney (2010)

[Itani, Kayssi and Chehab 2009] Itani, W., Kayssi, A. & Chehab, A. "Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures (DASC-09)"; Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing, IEEE, Chengdu, China (2009) 711-716

[Kaliski and Pauley 2010] Kaliski Jr, B. S. & Pauley, W. "Toward risk assessment as a service in cloud environments"; Proceedings of the 2nd USENIX conference on Hot topics in Cloud Computing, USENIX Association, Boston, MA (2010)

[Ko, Jeon and Morales 2011] Ko, S. Y., Jeon, K., & Morales, R. "The HybrEx Model for Confidentiality and Privacy in Cloud Computing"; Proceedings of the 2011 conference on Hot topics in Cloud Computing. USENIX Association, Portland, OR. (2011)

[Krautheim 2009] Krautheim, F. J. "Private virtual infrastructure for cloud computing"; Proceedings of the 2009 conference on Hot topics in Cloud Computing, USENIX Association, San Diego, CA (2009)

[Li et al. 2010] Li, A., Yang, X., Kandula, S. & Zhang, M. "CloudCmp: shopping for a cloud made easy"; Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing, USENIX Association, Boston, MA (2010)

[Mell and Grance 2010] Mell, P. & Grance, T. "NIST Definition of Cloud Computing v15". Washington, DC. (2010) Also appeared as an electronic version: http://www.csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.pdf

[Mowbray and Pearson 2009] Mowbray, M. & Pearson, S. "A client-based privacy manager for cloud computing"; Proceedings of the Fourth International ICST Conference on COMmunication System softWAre and middlewaRE (COMSWARE'09), ACM, Dublin (2009)

[Pal et al. 2011] Pal, S., Khatua, S., Chaki, N. & Sanyal, S. "A New Trusted and Collaborative Agent Based Approach for Ensuring Cloud Security"; Arxiv preprint arXiv:1108.4100 (2011)

[Peterson and Gondree 2011] Peterson, Z. & Gondree, M. "A Position Paper on Data Sovereignty: The Importance of Geolocating Data in the Cloud"; Proceedings of the 2011 conference on Hot topics in Cloud Computing. USENIX Association, Portland, OR. (2011) 1-5

[Popa et al. 2010] Popa, L., Yu, M., Ko, S. Y., Ratnasamy, S. & Stoica, I. "CloudPolice: taking access control out of the network"; Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks ACM, Monterey, CA (2010) 1-6

[Siebenhaar et al. 2011] Siebenhaar, M., Tsai, H. Y., Lampe, U. & Steinmetz, R. "Analyzing and Modeling Security Aspects of Cloud-based Systems"; GI/ITG KuVS Fachgespräch "Sicherheit für Cloud Computing", April (2011)

[Susman and Evered 1987] Susman, G. I. & Evered, R. D. "An Assessment of the Scientific Merits of Action Research"; Administrative Science Quarterly, 23, 4 (1978) 582-603

[Takabi, Joshi and Ahn 2010] Takabi, H., Joshi, J. B. D. & Ahn, G. "Security and privacy challenges in cloud computing environments". Security & Privacy, 8, 6 (2010) 24–31

[Tiwana et al. 2010] Tiwana, B., Balakrishnan, M., Aguilera, M. K., Ballani, H. & Mao, Z. M. "Location, location, location!: modeling data proximity in the cloud"; Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks, ACM, Monterey, CA (2010)

[Troncoso-Pastoriza and Pérez-González 2010] Troncoso-Pastoriza, J. R., & Pérez-González, F. "CryptoDSPs for Cloud Privacy"; Workshop on Cloud Information System Engineering (CISE'10) Hong Kong (2010), 1-12

[Wang, Chen and Zhou 2010] Wang, Chen & Zhou, Y. "A Collaborative Monitoring Mechanism for Making a Multitenant Platform Accountable"; Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing, USENIX Association, Boston, MA (2010), 18-25.

[Wang et al. 2009] Wang, C., Wang, Q., Ren, K., & Lou, W. "Ensuring data storage security in cloud computing."; Proceedings of the 17th International Workshop on Quality of Service (IWQoS), IEEE, Charleston, NC (2009), 1–9.

[Williams et al. 2011] Williams, D., Elnikety, E., Eldehiry, M., Jamjoom, H., Huang, H. & Weatherspoon, H. "Unshackle the Cloud!"; Proceedings of the 2011 Conference on Hot topics in Cloud Computing. USENIX Association, Portland, OR. (2011)

[Wood et al. 2010] Wood, T., Cecchet, E., Ramakrishnan, K., Shenoy, P., Van Der Merwe, J. & Venkataramani, A. "Disaster recovery as a cloud service: Economic benefits & deployment challenges"; Proceedings of the 2nd USENIX conference on Hot topics in Cloud Computing, USENIX Association, Boston, MA (2010), 1–7

[Yin 2009] Yin, R. K. Case Study Research (4th ed.). Sage, Inc., Thousands Oaks, CA (2009)