The SeCA Model

Thijs Baars, Utrecht University, The Netherlands

Marco Spruit, Utrecht University, The Netherlands

ABSTRACT

Security issues are paramount when considering adoption of any cloud technology. This chapter outlines the Secure Cloud Architecture (SeCA) model on the basis of data classifications, which defines a properly secure cloud architecture by testing the cloud environment on eight attributes. The SeCA model is developed using a literature review and a delphi study with seventeen experts, consisting of three rounds. We integrate the CI3A —an extension on the CIA-triad— to create a basic framework for testing the classification inputted. The data classification is then tested on regional, geo-spatial, delivery, deployment, governance & compliance, network, premise attributes. After this testing has been executed, a specification for a secure cloud architecture is outputted. The SeCA model is detailed with two example cases on the usage of the model in practice.

INTRODUCTION

According to both commercial reports as academic research, security issues are paramount when adoption of cloud solutions are being considered (Foster, Zhao, Raicu, & Lu, 2008; Ghinste, 2010; Mowbray & Pearson, 2009). However, no clear model exists to determine security issues and solutions.

Better yet, there is much debate which security threats and risks are applicable to computer networks, end-users or are actually cloud specific (Chen, Paxson, & Katz, 2010:4). They state that "arguably many of the incidents described as 'cloud security' in fact just reflect traditional web application and data-hosting problems [..] such as phishing, downtime, data loss, password weaknesses, and compromised hosts running botnets.". Moreover, they hold that most cloud security issues aren't new, but do need new implementations to provide the level of security wanted.

Therefore this chapter will provide an overview of the security issues and describe the Secure Cloud Architecture (SeCA) model to determine the security issues one might expect in a certain cloud environment and what solutions might be used to secure those issues. This framework will be developed by answering the following question:

Can the Cloud be a safe alternative for the storage and execution of organizational confidential data?

This model was developed in three steps. First, a literature review has been conducted. Second a delphi study was conducted to identify the perceived security issues by experts in the field. Third, the model was verified by the same experts in the last round of the delphi study.

By reading the overall themes in security, followed by cloud specific topics, an overview has been created that is used as the starting point in the development of the SeCA model.

A delphi study has been considered to be the best method for research in this chapter, as it provides the researchers with a qualitative data set which would allow to create and verify the model. It also allows the experts to see answers and be able to respond to these answers in upcoming rounds (Dalkey & Helmer, 1963). The first answer in question two is not per se answered by the same expert as answer one in question one, creating double-blind survey results. This way, a consensus can be reached on the various topics discussed in the delphi study. The delphi method was executed consisting of three rounds of surveys with qualitative questions. Three rounds were chosen instead of two, which is more common (Skulmoski, Hartman, & Krahn, 2007), so that a first round could be used to obtain general information on the topic, not specifically regarding to the model to be developed, while still having enough rounds to reach a consensus. The first round consisted of open questions where the experts were questioned on their experience with security and the cloud, issues and concerns regarding security in the cloud. These questions gave a wide result set that strengthened the results of the literature research earlier performed. Seventeen respondents answered all the questions in the survey in all three rounds, a rate of 65%. See Table 1 for an overview.

This model has been verified by an expert panel. The experts were selected on the basis of their function, publications and knowledge of security, the cloud or a combination thereof. This group of experts, 26 in total, comes from organizations within the business to consumer, business to business, business to government industries and governmental organizations. They were interviewed using the Delphi tool developed at the Wharton Business School (Wharton Business School, 2005).

#	**Position/function**	**Organization type**	**Cloud**	**Non-cloud**	**Security**
1	Consultant	Enterprise integrator	X	X	
2	Director	IT consultancy		X	X
3	Security consultant/architect	IT security firm		X	X
4	Researcher	American University,	X	X	X
5	Enterprise Architect	Enterprise transportation		X	
6	Sr. manager	Large accounting firm	X	X	X

7	Security advisor	Transportation firm		X	X
8	IT Architect	IT consultancy	X	X	
9	Manager	Security solutions	X		X
10	Security manager	Utilities		X	X
11	Consultant	IT consultancy	X	X	
12	Security manager	Government		X	X
13	Security manager	Healthcare products		X	X
14	Manager	IT consultancy		X	X
15	Consultant	Enterprise integrator		X	X
16	Security manager	Utilities		X	X
17	IT auditor	Accounting	X	X	X

Table 1: The experts (filtered on those that did all three rounds) in the delphi study

The burn chart below (table 2) shows the amount of consensus reached in the study, per topic addressed. White cells represent no questions in that topic were asked in that round; grey consensus was reached; chequered pattern a consensus in part was reached; black no consensus reached.

As one can see, not all topics reached consensus. This was due to the fact that in the expert selection business knowledge or technical knowledge on some topics were not taken into account. For example, the field of encryption is a very technical field that can be hard to fully understand and apply. Although some answers were very useful, other answers were dismissed in the same round as unfeasible, simpleminded or simply not true. This meant that the experience or knowledge between the experts varied too greatly to reach consensus. Subsequent research was done through literature review on the applicable topics.

Topic	**Round1**	**Round2**	**Round3**	**Consensus**	**Comments**
Security issues					All issues are accounted for in the model
Locationlessness					Location is a new issue and thoroughly discussed.
Trust issues					Outsourcing/insourcing/cloud differences are in discord

Encryption					Different knowledge levels; study done through literature.
Feasibility					Not a technical/security issue. Topic abandoned.
Model					Model validated and approved by the experts.
Auditing					Issues reached consensus; added to the CI3A.

Table 2: burn chart of the consensus reached in the delphi study.

BACKGROUND

The Cloud thus far has attracted a lot of attention, this section tries to provide a concise overview of key research and models previously developed.

The Jericho Forum (2009) has previously modeled the cloud in order to help users understand the different facets of the cloud and support a secure use of cloud technologies. It does however not account for the complexities seen in cloud security. Siebenhaar, Tsai, Lampe, & Steinmetz (2011) describe a holistic model for analyzing and modeling security aspects of cloud-based systems, Almorsy, Grundy, & Ibrahim (2011) provide a collaboration based framework for determining security management. The presented SeCA model in this research differs from the previously mentioned as it focuses on security measures within the architecture of the cloud determined by data classifications.

Subashini & Kavitha, (2011) discuss a survey concerning security issues in specific delivery models of the cloud. Although the encompassing research should lead to a model, that model is still under development. Benson, Sahu, Akella, & Shaikh (2010) discussed security issues from a cloud provider's support division perspective by looking at SLA structures. They report the most commonly found problems in IaaS architectures and offer three practical solutions. None of these are technical security solutions. Kaliski Jr & Pauley (2010) discuss risk assessment of the cloud, stating that "[t]he very characteristics that make cloud computing attractive also tend to make it hard to assess" (p.2). Richter et al. (2011) discuss VM retrospection, a method for inspecting previous VM states in order to perform forensics, debugging, troubleshooting and so forth. Christodorescu, Sailer, Schales, Sgandurra, & Zamboni (2009) discuss methods of securing Clouds at the Virtual Machine (VM) level. They provide a short overview of known VM issues and solutions, and then propose their system which protects VMs in a cloud against malware and rootkits using a white/blacklist method. Wang, Wang, Li, Ren, & Lou (2009) discuss the necessity of a Third Party Actor (TPA) to assure security standards and to provide transparency in the security controls. Jensen, Schwenk, Gruschka, & Iacono (2009) describe technical security issues related to cloud, using the Amazon EC^2 cloud as a case. Although all issues discussed are related to the cloud, all them were already apparent before the coming of the cloud, some just

have found new grounds to be relevant again. Vigfusson & Chockler (2010) discuss in "Clouds at Crossroads: Research Perspectives" research topics in the cloud, including privacy related issues. Discussing the trust problems that arise with the complex models of some cloud environments, it provides a few solutions to suggestions that might be very feasible.

The research presented in this chapter overlaps with current research in that it provides an insight in cloud security, introducing and explaining it, but it also expands the current research with proposed solutions to solve cloud security issues at the managerial level not discussed so far. The current research explains either very technical details on protecting the cloud, where the mere describes arbitrary issues that are not specific to the cloud, or gives overviews of the cloud where security issues are touched lightly, this research will focus on the security issues in depth that come with the cloud in a more practical sense. We specify cloud specific issues and general security issues that have found new terrain in the cloud environment. To conclude, this research provides users with a model that navigates them through the security checkpoints in cloud environment, outputting an architecture specific for their data classification.

RISKS & THREATS IN THE CLOUD

This section describes the main overview of the security issues, risks and threats in the cloud, which are all part of the SeCA model which is presented on page 11. First, the CI3A is described, followed by all the attributes in the SeCA model: regional, geo-spatial, delivery model, deployment model, governance & compliance, network and premise.

CI3A

Because of the complexities the cloud presents as dictated by a majority of the participants of the delphi study, the *de facto* CIA triad (*e.g.* Avizienis, Laprie, Randell, Landwehr (2004)), which is used for testing the confidentiality, integrity and availability in systems, data flows and so forth, was found to be too constrained. For that reason, we developed the CI3A as an extension of the CIA triad, comprising of confidentiality, integrity, availability, accountability and auditability. See Figure 1. The proposed model utilizes CI3A to assure the right level of security is maintained within the environment. This section will describe the CI3A, following separate sections on locationlessness and trust chains.

{INSERT: CI3A.tif}

Figure 1: The CI3A visualized

Confidentiality is reached by proper authentication/authorization controls and encryption methods such as secured computing and two-factor encryption. Preventing data leakage is a central part within the confidentiality strategy. The choice of distribution and delivery model influences the level of confidentiality and the methods needed to assure confidentiality.

Integrity assures only authorized actors have access to certain data and that data gets distributed to only authorized persons. Within that distribution, any editing or changes within the data should only be made by the right persons. Governance and Compliance influence the integrity of

the data; a fully compliant environment is more likely to assure integrity. As with Confidentiality, the chosen delivery and distribution model influences the level of integrity.

Availability comprises of measures to prevent unauthorized actors from deleting and moving data, or accessing those files, minimizing downtime of the environment. These measures could be a HA infrastructure, strong authentication servers, disaster recovery or external hosted service such as CloudPolice (Popa, Yu, Ko, Ratnasamy, & Stoica, 2010). Availability plays a big role within the cloud environment, as servers can be hosted anywhere in the world, at multiple locations. Although an advantage in the eyes of HA and disaster recovery; latency, desynchronization and vulnerabilities in the transceiver links can pose threats. Also, ownership of data is a part of availability. Availability is linked to regional, geo-spatial, network, premise and to the delivery and distribution models.

Accountability defines the measures taken to assure that no actor can make actions without a record. This is needed for forensics and governance. The measures needed to assure accountability greatly depend on the delivery model, but also on the distribution model and compliance in general. (Chen Wang & Zhou, 2010) have found accountability of paramount importance in the cloud, proposing a method for transferring accountability onto an external host in order to perform accountability in a multitenant platform.

Auditability, the ability of the environment to be audited, is directly related to governance and compliancy. Without a decent grade of auditability, compliance cannot be achieved. Auditability is influenced by the delivery and distribution model, as with the geo-spatial and geographic boundaries.

The Regional Attribute

Regional describe the boundaries that signify separate legal systems. These boundaries include cities, states, countries and territories. Some changes in legal systems might be significant, such as the difference in respect to privacy between the European Union and China; some might be incremental such as the difference between county and state laws in the United States. These differences however do impose a risk if your data gets placed on a physical server crossing such a boundary. Next to that, different legal systems have different perspectives on privacy, the use of subpoenas on data extraction from datacenters. As one expert commented: “bringing privacy information out of the European Union can be [a] violation of local or European law”. This would mean that keeping in compliance with laws, be it local, national or international, will become more difficult without knowledge of the physical location of the data store and computing unit. Peterson & Gondree (2011) provide an elaborate view on the importance of data location awareness from an American perspective.

The Geo-spatial Attribute

With geo-spatial risks, the distance of objects “relating to the relative position [..] on the earth's surface” (Collins English Dictionary, 2009) is meant, in this case the distance between servers, but also the location of each server. This can be of importance in the case of disaster recovery, but also with regards to physical security as presented in security norms such as the ISO 2700x series. In the light of location, one could also consider other features such as the building type,

the accessibility of the server etc. Geographic location should also be taken into account in the light of latency and propagation speeds, as emphasized by Tiwana, Balakrishnan, Aguilera, Ballani, & Mao (2010).

The Delivery Model

The cloud has three distinct platforms on which a cloud environment can be offered. They are stackable, meaning that if you have a Software as a Service (SaaS) solution, chances are that your provider manages a Platform as a Service (PaaS), but takes services from an Infrastructure as a Service (IaaS) provider. See also Figure 2. This, however, does not mean that every SaaS solution is running as the top of a stack of cloud platforms. A SaaS solution can run on a traditional hardware stack with no further cloud environment attached.

{INSERT: trust_chains.tif}

Figure 2: trust chains in cloud architectures

Trust is a major issue in any relation, be it personal or professional. Although this is trivial, cloud computing can create trust chains, in which the end user is not always aware which other links are present in his chain of trust. This pertains especially towards delivery models. With IaaS, the tenant is in direct contact with the owner of the infrastructure (in some cases there might be a reseller in between) who can have outsourced duties associated with the maintenance of the physical systems. In a SaaS model, one is not aware if the SaaS provider also owns the platform, or the infrastructure. This means that there might be a variety of different actors working on the cloud, whom all might be able to access the data that is being used in the SaaS in some way or another. Actors whom the tenant initially didn't trust have now become a part of his organizational network. This might result in actions that are a threat to the data. Although doing business is about making relations and trust, yet not insurmountable, they are a risk factor.

The Deployment model

The cloud comes in four different deployment models, these are private, public, hybrid and community/Partner clouds. The difference between these four models is the openness of the cloud to its tenants.

In a private cloud, the cloud infrastructure is operated for just one organization. This does not mean that it has to be managed by that organization. The management of the private cloud can be done by a third party, and the cloud itself can be physically located on the premises of that organization, or can be hosted somewhere else (sometimes called a virtual private cloud.) The cloud can exist behind a firewall of the organization, and thus only accessible within its private network, but can also be hosted off-premise on dedicated hardware (thus no multitenancy with other organizations). These are all factors that influence the security risks. The main difference between a mainframe or internal traditional datacenter and a private cloud is that there is a virtualization layer that can be used to host SaaS applications, rapid deployment and other benefits of cloud computing. Even though private clouds conventionally lack the flexibility of their public equivalents, a model has been proposed to allocate public cloud space for private clouds, giving it the full flexibility as public clouds with the added security of private clouds (Ko, Jeon, & Morales, 2011).

In a community cloud, a community or group of organizations share the same cloud infrastructure. These communities have shared concerns, such as a mission, goal and/or policy. The cloud can be managed by one of the organizations within the community or by a third party, and may exist on or off premise (Mulholland, Pyke, & Fingar, 2010).

In a Public Cloud, the cloud is open for use to a large group of tenants, which do not need to know each other. The cloud is a ran by a cloud service provider. An example can be a majority of offerings from Force.com and Google's Gmail to VPS.net, Rackspace cloud hosting and other public services, free or on a subscription basis. Typically these are off premise, out of the organizational network.

A Hybrid Cloud is a composition, or hybrid if you will, of two or more clouds of the types mentioned above. They are unique entities, but tied together with APIs to enable the exchange of data and applications. Due to the nature of the different clouds the hybrid consists of, it can be deployed both on and of premise, and be fully, partly or not behind the firewall of the organization.

Governance & Compliance

Executing governance and compliance is according to our experts is a much debated issue. Because governance and compliance greatly depend on the infrastructure of the system and the above mentioned boundary issues, this topic is much under the discretion of the chosen cloud environment.

Depending on the chosen delivery model, compliance can be completely out of hand. A SaaS application depends on their vendors for governance and compliance. For PaaS it is partly the same, any compliance and governance within the software and how it handles data is on the part of the developer. The governance of the infrastructure and platform on which the application relies is in the hands of the provider. As with SaaS, negotiations need to take place with the provider in order to secure compliance. For IaaS, most of the governance and compliance lays in the hands of the tenant. The IaaS provider has to take care of the compliance to standards such as SAS-70, but many issues like privacy, data encryption and authentication are the responsibility of the tenant.

Concerning deployment models, compliance and governance in a public cloud can be difficult, as you are limited to your VM instance, whereas in a private cloud your negotiation position will be stronger as there are no other tenants to take into account. In a partner cloud, one can imagine that governance and compliance is a shared goal.

Concerning boundaries, the major aspect is the geographic location of the servers. The easiest option is of course in the same region as the organization resides, most knowledge of laws and executing governance/assuring compliance will be readily available. Auditing will not be an issue, as you can identify an auditing partner with whom you can easily communicate. That being said, the hardest option is obviously a cloud environment dispersed over the globe. Although disaster recovery wise there will be no issue complying to the toughest guidelines, getting audited and governance worldwide will be tougher. Although experts in the survey were

wary of the fact that it could be done, in a personal interview with an Chief Information Security Officer of a large utilities company, it was made clear that a global audit is unprecedented.

Network

Network indicates the boundary of an organizational computer network. This is an important factor, as some information is not wanted outside the corporate network, such as trade secrets. Keeping a cloud environment within the boundaries of the network can be reached by keeping it on premise and thus physically in the network, or it can be reached by creating a VPN connection (as elaborated by Wood, Gerber, Ramakrishnan, Shenoy and Van der Merwe (2009)) or a VLAN (in case of internal networks, or in public as described by Hao, Lakshman, Mukherjee and Song (2010)) in order to keep the information within the network.

Because some configurations stretch the extension of the enterprise network, additional risks are incurred due to this stretch, as some of our experts mentioned in the survey. This stretch in the network is also noticeable in the added amount of actors which have to be trusted. The cloud provider will probably have access to your network, or the possibility to illegally gain so.

An added risk is the uncertainty of the WAN infrastructure at the providers side. Connecting with the cloud provider might create vulnerabilities that could threaten the corporate network. Next to that, multitenancy might also be considered within the range of network boundaries. Although multitenancy should never be a threat to the virtual machine, in that it shouldn't have the possibility of other tenants to enter your VM, it has been proven that a vulnerability on the OS level could provide access to other VMs. Ristenpart, Tromer, Shacham and Savage (2009) describe ways to discover where nodes are hosted on Amazon's EC2 cloud, following with a discussion how to place a co-resident on that physical server in order to able to reach the hardware a selected node is on. By then compromising the system, the selected node might be entered. This is a risk that has to be considered, how small it seems to be (see (Asadoorian, 2007; Mehta & Smith, 2007; Ormandy, 2007) for an overview).

On or Off Premise

Organizational premises play a role in the physical location of the cloud environment. One can either wise choose to have the hardware reside on or off organizational premises. For high security purposes keeping the hardware on premise, and thus fully in one's control, might provide a benefit; personnel can be screened, there's extended access control to the datacenter and forensics. This extends the discussion on the geographic location of the server, presenting a trade off in security between on-premise servers versus geo-spatial choices.

ENCRYPTION IN THE CLOUD

Encryption plays a vital role within the cloud environment. It is affected by all but the geo-spatial attributes in the SeCA model and affects the regional, delivery and deployment model. Although encryption is a broad topic that has been covered in many papers, theses and books, there are some aspects that are specifically related to the cloud. VPN tunnels, together with SSH can provide secure access to the cloud environment. Two-factor can be very helpful for the cloud environment. Many institutions are using hardware key-tokens or SMS gateways in order two

provide the second form of authentication apart from keying in a password. Authentication servers using protocols as RADIUS in combination with LDAP, Kerberos or Active Directory can handle all access requests in a proven manner as they are no different from any LAN/WAN setup at a traditional environment. The author therefore believes that in terms of access control, authentication and authorization, no cloud specific issues are at hand.

Apart from the aforementioned, an encryption method specifically pertaining to the cloud is secure computing. Secure computing offers a solution to issues that arise when multiple systems have to use secure information in transactions and computations, in essence described by Yao's (1982) Millionaires' problem. This research has been extended by Goldreich (2000), who researched the problem with multiple actors called Secure Multi-party Computations (SMCs). Recent research involves SMC geometry, researching transactions of polygons on convex hulls. See (Wang, Luo, & Huang, 2008) for an overview.

It is known that any multi-party computational problem can be solved using the generic technique of Yao (Yao, 1982). To overcome the overhead with Yoa's Millionaires' problem, and thus SMC, it seems that algorithms designed to compute a special task need to be written (Feigenbaum, Pinkas, Ryger, & Saint Jean, 2004; Goldreich, 2000). Using encryption methods such as homomorphic encryption and public key encryption, several algorithms have shown to be applicable to the cloud (Das & Srinathan, 2007; Hu & Xu, 2009; Troncoso-Pastoriza & Pérez-González, 2010) and have proven to provide the security needed for the cloud within test situations approaching real life cloud environments.

These methods of secure computing would allow the creation of a chain of trust that is secure, even though not all parties within the chain know each other nor trust each other. This could overcome any trust issues that might be in the field of cloud environments. Together with the enhanced and proven techniques of authentication and authorization already available, encryption can make the cloud a very secure architecture.

Apart from the above mentioned, Mowbray & Pearson (2009) have developed a privacy manager that can obfuscate data in effort to protect it from malicious providers.

The following table shows how encryption affects and is affected by the choice of certain attributes in the model. For example, when analyzing the attribute Compliance in the model the encryption attribute is also influenced, as some standards and classifications define levels of encryption needed. Defining the attribute Encryption influences the CI3A on all but availability as the means of encryption can assure the confidentiality, integrity, accountability and auditability of an architecture.

SeCA attribute	Confidentiality	Integrity	Availability	Accountability	Auditability
Regional	X			X	X
Geo-spatial			X		
Compliance	X	X			X
Delivery model	X	X	X		X
Deployment model	X	X	X	X	X

Encryption	X	X		X	X
Network	X		X	X	
Premises	X	X	X		

Table 3: Encryption and how it is affected or affects the other attributes in the SeCA model

THE SECA MODEL

Resulting from the research conducted, we can summarize that the cloud can be secure, as long as its policies and SLAs are correctly in place and enforced. The different factors and risks involving cloud computing make it difficult to pinpoint to one secure cloud. In fact that is impossible, due to the diversity of cloud architectures and the data that is being stored on it. To circumvent this problem, the SeCA model has been designed. This model has been validated in the final round of the delphi study.

The model described below, the SeCA Model, gives an abstract overview of all the characteristics of the cloud. Figure 3 depicts the SeCA model. It defines a secure cloud architecture for a specific data classification.

{INSERT: seca_model.tif}

Figure 3: The SeCA Model

The model outputs guidelines for the cloud environment and to which specification a cloud solution should adhere. 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. It is assumed that there are data classifications defined within the organization.

The organization in this example is trying to identify whether their standard office suite can be replaced by a cloud offering. They have identified two cloud service providers that offer office suites, one by the ACME corporation and one by WEC, Inc. They are analyzed in in table 5. We will refer to this matrix in both sections. It depicts a matrix which describes the values of the attributes in the SeCA model for both cloud solutions.

To keep this example brief, we will only discuss a fictional data classification named “private”. This classification is shown in table 4. This classification can be created by using the template shown below (template 1.) The template allows for one to map an existing data classification to the SeCA model, providing a clear overview of the attributes and their values. The final classification as shown in table 4 shows the SeCA attributes and the corresponding values. It is already mapped to the SeCA model. One can imagine that the classification before mapping does not literally mention the attributes, nor talk about deployment and delivery models. Some attributes, like encryption, might be already present in the classification before being mapped to the SeCA model.

Classification Name/ identification: _____ Expert's Name:_____

Regional:	
Geo-spatial:	
Governance & Compliance:	
Delivery Model:	○ IaaS ○ PaaS ○ SaaS
Deployment model:	○ Private ○ Partner/Community ○ Public ○ Hybrid
Encryption:	
Network:	○ Within ○ Outside ○ Any
Premises:	○ On premise ○ Off premise ○ Any

Template 1: SeCA Data Classification Template

Classification: Private	**Secure Cloud Architecture specification**
Attribute	**Value**
Regional	Cloud environment physically within the same region as the organization.
Geo-spatial	A HA architecture is preferred, at least one backup location in a different building on separate power net.
Governance/compliance	No need to adhere to specific standards. Annual audit is required to assure proper protection.

Delivery model	Any
Deployment model	Any
Encryption	A proper authentication and authorization system should be in place for any actions. Alignment with the LDAP server in place is preferred.
Network	Any
Premises	Any

Table 4: results from the cloud analysis

Supplier	**Service**	**Deployment Model**	**Delivery Model**	**Location**		
				jurisdiction	**premise**	**Network**
WEC	Office Suite	Public	SaaS	Safe harbor, World wide	Off	Outside
ACME	Office Suite	Public	SaaS	USA	Off	Outside, inside through VPN access

Table 5: Service Matrix

Security				
backup	**location**	**A&A**	**Encryption**	**Certification**
Synchronization, mirroring and disaster recovery	Worldwide (no particulars known)	LDAP, proprietary user/pass authentication	SSL, TLS. No storage encryption	SAS 70-II, SOX, PCI-DSS, Safe Harbor
Disaster recovery Location in the EU, Mirroring within the USA.	USA	AD setup, custom two-factor auth.	256bit SSL, TLS	ISO 27001, safe harbor, SAS 70-II.

Table 5: Service Matrix -continued

CI3A				
confidentiality	**Integrity**	**accountability**	**auditability**	**availability**
No dedicated hardware possible	SSL, third-party apps	Full logging, data always belongs to the end user, 3rd party apps	Global internal audits, no specifics available	99.9% uptime, scheduled downtime
Possible to get dedicated hardware	SSL only	End user is responsible	unclear	99.99% uptime

Table 5: Service Matrix -continued

Forward Direction

When the model is used with a forward direction, the assessor at the organization inputs a data classification. From this data classification a set of rules is created to which the cloud architecture should adhere. A set of outer-boundaries if you will. These are shown in table 4.

Once this is done, a list of cloud providers/solutions who can adhere to the results from the assessment is created. This can be done manually or by using proposed services such as CloudCMP (A. Li, Yang, Kandula, & Zhang, 2010). Ultimately, a cloud provider is selected, arrangements are made and the data can be placed in the cloud.

The flow chart below (Figure 4) shows where this assessment ordinarily should take place.

{INSERT: process_forward.tif}

Figure 4: The position of the SeCA model of transitioning data into the cloud in a forward direction (simplified)

It can occur that each classification has a different output from the assessment. (it is actually most likely to do so.) In that case several options are open. For each classification a different list of cloud providers is made in order to find and select the right cloud provider who can provide the cloud architecture needed. These can be combined in Hybrid Clouds. One can also decide that for certain classifications it is simply not feasible to transfer that data into the cloud and thus stay with the solutions already in place.

The model does not provide the intelligence which classifications could be hosted at the same cloud architectures. It is for the assessor to decide which cloud architectures that are the result from the assessments can be merged.

This classification will normally be created by a security officer in the organization. He can use the SeCA Model to describe the values for the attributes in the model. On this basis the lower and upper bound values for the architecture are drafted. One can then look at the services that fit within these bounds and deem them safe enough, or not, for the data that'll be used with that service.

Backward Direction

A different way of using the model, the so-called backward direction, is to select the services one is interested in. By investigating the properties of the service and its underlying architecture a matrix can be created. This matrix, shown in table 5 with two fictional office services, can then be mapped to a data classification, shown in table 4, specified by the organization. As we can see from the examples given, the Office Suite service from the ACME Corporation cannot be used for data classified as "private", as the classification clearly states that an "Annual audit is required to assure proper protection." In the Matrix, it states that in the column of the CI3A→auditability: "unclear".

This method can be easier if one wants to know if a certain cloud solution could security-wise work within the organization.

As one can see, different cloud solutions have different attribute values. It is for this reason that this direction can be preferred as it provides a great overview of the services that are being considered. Even if no effort has been put into mapping the classifications to the SeCA model, one can already get a hunch on the solution, how the architecture is setup and where the data is located.

Figure 5 provides a simplified flowchart of the backward direction. The differences with the forward direction (figure 4) are clear: one starts with listing the cloud solutions instead of doing that at a later stage, after which the solutions are mapped to data classifications. Therefore, this method provides an early overview of which solutions fit with which data classification.

{INSERT: process_backward.tif}

Figure 5: The position of the SeCA model of transitioning data into the cloud in a backward direction (simplified)

FUTURE RESEARCH DIRECTIONS

Further research can be conducted in the legal field. This was out of scope of this research, but the legal issues surrounding auditing, SLAs and NDAs are of paramount importance for the security in the cloud. SLAs especially, are of profound importance as they describe what measures a cloud provider should undertake for the security of the cloud. This paper unfortunately has not had the possibility to explore the provider side of the cloud environment much.

Related to this is auditing in international/worldwide clouds. Auditing certifications, governance and compliance to legal systems in these environments means that auditing firms, datacenter owner, providers and application owner all need to work together in order to get successful audit. In international and worldwide clouds these relations might become very complex, not to mention that multiple audit firms/offices have to work together. The issues raised with datacenters situated in different legal regions, such as China and the United States, are worth more research. Auditing plays also here a major role.

A pressing issue not discussed, but worth the research are third party appliances that are currently installed in traditional datacenters. These appliances cannot be directly converted to the cloud, as the cloud does not offer any place for such appliances. It seems that at the moment of writing many of these appliances are converted to the cloud by their developers. It is nonetheless interesting to see what impact these appliances have on the adoption of the cloud. On that note, (Krautheim, 2009) has developed an infrastructure called PVI for the cloud that automates provisions depending on security settings. It would be interesting to see how the SeCA model can be connected to the presented PVI.

Although some cloud providers are certified, the impact of that certification on the real security of the services the provider offers is not always known. SAS70 for example does not offer any concrete security, it only offers a framework for auditing internal controls. The cloud provider will need to list its internal controls for any user to see what has been audited. It might be

interesting to see how cloud providers use that information, what they do with it and whether the certifications really add up to extra level of security that is said it adds.

A case study should be conducted putting the model to practice, validating it in the work field. The validation so far has been merely theoretical by the experts in the delphi study.

CONCLUSION

Defining something as secure depends on many factors. Depending on the sort of data, the classification of that data and taking that wholly into perspective of the cloud environment, it can be said that the cloud is secure in certain situations. Depending on the outcomes of investigations, there should always be a cloud architecture that fits one's security needs. Better yet, the cloud can provide additional layers of security by utilizing virtualization, elasticity and HA architectures. Even though the additional layer of virtualization on the system might provide additional hazards, looking at the scarcity of exploitations in this layer one can rationally say that the virtualization layer adds more protection than threats.

By using the SeCA model described above, each and every classification can be checked to see how a cloud architecture should be designed in order to meet the security standards needed. It will, however, depend on the cloud provider whether it can deliver the architecture that is needed.

For the upmost secure classifications, a private cloud, hosted on premise, within the network, with mirroring on a different physical location (branch office) utilizing the needed encryption methods will provide a very secure architecture whilst maintaining some of the flexibility the cloud has to offer.

For every architecture counts that data location awareness is essential. Without the full knowledge of where the data resides and is processed, issues will arise in all actors of the CI3A. Data location awareness will also provide the means for compliance, legally and to security standards. These standards are being adopted by all major vendors, including Amazon, Google and Microsoft, with smaller ones following. This facilitates full compliance to the de facto security and auditing standards such as SAS 70, ISO 27000 series, PCI and COBIT. It depends, once again, on the configuration of the cloud architecture and where applicable the willingness of cloud provider to allow for audits. If the selected cloud architecture features datacenters in widely spread different parts of the world, auditing might be more complicated. This of course also applies to the compliance to legal systems (privacy, intellectual property and auditing regulations) which can vary between jurisdictions. It is because of these implications that so-called locationless clouds are not preferable. They have an opaque layer that hides the user from vital knowledge in order to gain assurance from the CI3A.

The model presented covers the complexity of the cloud and gives implementers, decision makers and IT professionals a hands-on tool for deciding whether and which cloud solution to use for their data.

REFERENCES

Almorsy, M., Grundy, J., & Ibrahim, A. S. (2011). Collaboration-Based Cloud Computing Security Management Framework. *Cloud Computing (CLOUD), 2011 IEEE International Conference on* (pp. 364–371). Washington, DC.: IEEE. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6008731

Asadoorian, P. (2007). Escaping from the Virtualization Cave. *PaulDotCom*. Retrieved July 12, 2011, from http://www.pauldotcom.com/2007/07/31/escaping_from_the_virtualizati.html

Avizienis, A., Laprie, J., Randell, B., & Landwehr, C. (2004). Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing, 1*(1), 11–33.

Benson, T., Sahu, S., Akella, A., & Shaikh, A. (2010). A first look at problems in the cloud. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (pp. 1-7). Boston, Massachusetts: USENIX Association. doi:10.1.1.150.2883

Chen, Y., Paxson, V., & Katz, R. H. (2010). *What's new about cloud computing security*. *University of California, Berkeley Report No. UCB/EECS-2010-5 January* (Vol. 20). Berkeley, CA. Retrieved from http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-5.pdf

Christodorescu, M., Sailer, R., Schales, D. L., Sgandurra, D., & Zamboni, D. (2009). Cloud security is not (just) virtualization security. *Proceedings of the 2009 ACM workshop on Cloud computing security - CCSW '09* (p. 97). New York, New York, USA: ACM Press. doi:10.1145/1655008.1655022

Dalkey, N., & Helmer, O. (1963). An Experimental Application of the DELPHI Method to the Use of Experts. *Management Science, 9*(3), 458-467. JSTOR. doi:10.1287/mnsc.9.3.458

Das, A. S., & Srinathan, K. (2007). Privacy Preserving Cooperative Clustering Service. *15th International Conference on Advanced Computing and Communications (ADCOM 2007)* (pp. 435-440). Guwahati, India: Ieee. doi:10.1109/ADCOM.2007.52

Feigenbaum, J., Pinkas, B., Ryger, R. S., & Saint Jean, F. (2004). Secure computation of surveys. *EU Workshop on Secure* (pp. 1-6). Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.100.3258&rep=rep1&type=pdf

Forum, J. (2009). *Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration*. Retrieved from http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

Foster, I., Zhao, Y., Raicu, I., & Lu, S. (2008). Cloud computing and grid computing 360-degree compared. *2008 Grid Computing Environments Workshop* (pp. 1-10). Austin, TX: Ieee. doi:10.1109/GCE.2008.4738445

Ghinste, B. V. (2010). Gartner: Private Cloud Computing Plans From Conference Polls. *MSDN Blogs*. Retrieved June 27, 2011, from http://blogs.msdn.com/b/architectsrule/archive/2010/05/07/gartner-private-cloud-computing-plans-from-conference-polls.aspx

Goldreich, O. (2000). Secure multi-party computation. *Working Draft*. New York, New York, USA: Citeseer. doi:10.1145/508172.508174

Hao, F., Lakshman, T., Mukherjee, S., & Song, H. (2010). Secure cloud computing with a virtualized network infrastructure. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (pp. 16–16). Boston, Massachusetts: USENIX Association. doi:10.1234/12345678

Hu, H., & Xu, J. (2009). Non-Exposure Location Anonymity. *2009 IEEE 25th International Conference on Data Engineering* (pp. 1120-1131). Shanghai, China: IEEE. doi:10.1109/ICDE.2009.106

Jensen, M., Schwenk, J., Gruschka, N., & Iacono, L. L. (2009). On Technical Security Issues in Cloud Computing. *2009 IEEE International Conference on Cloud Computing* (pp. 109-116). Bangalore, India: IEEE. doi:10.1109/CLOUD.2009.60

Kaliski Jr, B. S., & Pauley, W. (2010). Toward risk assessment as a service in cloud environments. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (pp. 13–13). Boston, Massachusetts: USENIX Association. Retrieved from http://portal.acm.org/citation.cfm?id=1863116

Ko, S. Y., Jeon, K., & Morales, R. (2011). The HybrEx Model for Confidentiality and Privacy in Cloud Computing. *Proceedings of the 2011 conference on Hot topics in cloud computing*. Portland, OR. Retrieved from http://www.usenix.org/event/hotcloud11/tech/final_files/Ko.pdf

Krautheim, F. J. (2009). Private virtual infrastructure for cloud computing. *Proceedings of the 2009 conference on Hot topics in cloud computing* (pp. 5–5). San Diego, CA: USENIX Association. Retrieved from http://portal.acm.org/citation.cfm?id=1855538

Li, A., Yang, X., Kandula, S., & Zhang, M. (2010). CloudCmp: shopping for a cloud made easy. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (pp. 5–5). Boston, Massachusetts: USENIX Association. Retrieved from http://portal.acm.org/citation.cfm?id=1863108

Mehta, N., & Smith, R. (2007). VMWare DHCP Server Remote Code Execution Vulnerabilities. *IBM Internal Security Systems*. Retrieved July 12, 2011, from http://www.iss.net/threats/275.html

Mowbray, M., & Pearson, S. (2009). A client-based privacy manager for cloud computing. *Proceedings of the Fourth International ICST Conference on COMmunication System softWAre and middlewaRE - COMSWARE '09* (p. 1). Dublin, Ireland: ACM Press. doi:10.1145/1621890.1621897

Mulholland, A., Pyke, J., & Fingar, P. (2010). *Enterprise Cloud Computing*. Tampa, FL: Meghan-Kiffer Press.

Ormandy, T. (2007). An empirical study into the security exposure to host of hostile virtualized environments. *Proceedings of CanSecWest Applied Security Conference*. Vancouver., Canada: Citeseer. doi:10.1.1.105.6943

Peterson, Z., & Gondree, M. (2011). 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*. Portland, OR. Retrieved from http://www.usenix.org/event/hotcloud11/tech/final_files/Peterson.pdf

Popa, L., Yu, M., Ko, S. Y., Ratnasamy, S., & Stoica, I. (2010). CloudPolice: taking access control out of the network. *Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks* (p. 7). Monterey, CA: ACM. Retrieved from http://portal.acm.org/citation.cfm?id=1868454

Richter, W., Ammons, G., Harkes, J., Goode, A., Bila, N., de Lara, E., Bala, V., et al. (2011). Privacy-Sensitive VM Retrospection. *Proceedings of the 2011 conference on Hot topics in cloud computing* (pp. 1-6). Portland, OR. Retrieved from http://www.usenix.org/events/hotcloud11/tech/final_files/Richter.pdf

Ristenpart, T., Tromer, E., Shacham, H., & Savage, S. (2009). Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. *Proceedings of the 16th ACM conference on Computer and communications security* (pp. 199–212). Chicago, IL: ACM. doi:10.1.1.150.681

School, W. B. (2005). Delphi Decission Aid. Retrieved October 5, 2010, from http://armstrong.wharton.upenn.edu/delphi2/

Siebenhaar, M., Tsai, H. Y., Lampe, U., & Steinmetz, R. (2011). Analyzing and Modeling Security Aspects of Cloud-based Systems. *GI/ITG KuVS Fachgespräch "Sicherheit für Cloud Computing"*, (April). Retrieved from ftp://ftp.kom.e-technik.tu-darmstadt.de/papers/STLS11.pdf

Skulmoski, G. J., Hartman, F. T., & Krahn, J. (2007). The Delphi Method for Graduate Research. *Journal of Information Technology Education, 6*, 1-21. doi:10.1.1.151.8144

Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, *34*(1), 1-11. Elsevier. doi:10.1016/j.jnca.2010.07.006

Tiwana, B., Balakrishnan, M., Aguilera, M. K., Ballani, H., & Mao, Z. M. (2010). Location, location, location!: modeling data proximity in the cloud. *Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks* (p. 15). Monterey, CA: ACM. doi:10.1145/1868447.1868462

Troncoso-Pastoriza, J. R., & Pérez-González, F. (2010). CryptoDSPs for Cloud Privacy. *Workshop on Cloud Information System Engineering (CISE'10)* (pp. 1-12). Hong Kong, China. doi:10.1.1.185.429

Vigfusson, Y., & Chockler, G. (2010). Clouds at the crossroads. *Crossroads*, *16*(3), 10-13. doi:10.1145/1734160.1734165

Wang, Chen, & Zhou, Y. (2010). A Collaborative Monitoring Mechanism for Making a Multitenant Platform Accountable. *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing* (pp. 18-25). Boston, Massachusetts: ACM. Retrieved from http://www.usenix.org/event/hotcloud10/tech/full_papers/WangC.pdf

Wang, Qian, Wang, C., Li, J., Ren, K., & Lou, W. (2009). Enabling Public Verifiability and Data Dynamics for Storage Security in Cloud Computing. In M. Backes & P. Ning (Eds.), *LNCS, ESORICS 2009* (5789th ed., Vol. 5789, pp. 355-370). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-642-04444-1

Wang, Qi, Luo, Y., & Huang, L. (2008). Privacy-preserving Protocols for Finding the Convex Hulls. *2008 Third International Conference on Availability, Reliability and Security*, (070412043), 727-732. Ieee. doi:10.1109/ARES.2008.11

Wood, T., Gerber, A., Ramakrishnan, K., Shenoy, P., & Van der Merwe, J. (2009). The case for enterprise-ready virtual private clouds. *Proceedings of the 2009 conference on Hot topics in cloud computing* (pp. 4–9). Monterey, CA: USENIX Association. Retrieved from http://portal.acm.org/citation.cfm?id=1855537

Yao, A. C. (1982). Protocols for secure computations. *23rd Annual Symposium on Foundations of Computer Science* (pp. 160-164). Chicago, IL: IEEE. doi:10.1109/SFCS.1982.38

ADDITIONAL READING

As topics quickly evolve and develop in this section, we believe that regular literature normally doesn't keep up. Therefore, take a look at the links provided below.

http://www.usenix.org/event/ The Usenix event calendar. The authors follow the Hot-ICE and HotCloud workshops, which provide with relevant information, but other workshops include interesting material as well. Work presented in previous events is available through their respective websites.

http://www.sigcomm.org/ The ACM special interest group for computer communication and computer networks. The hosted conferences and workshops (especially SIGCOMM conference and the Hotnets workshop) present very interesting material, as does the IEEE/ACM Transactions on Networking.

KEY TERMS

Cloud computing: a disruptive delivery model which turns the internet into a storage and computing power provider.

CI3A: an extension of the traditional CIA triad which models confidentiality, integrity, availability, accountability and auditability aspects.

SeCA: the Secure Cloud Architecture (SeCA) model which helps determine the security issues one can expect in a certain cloud environment based on data classification through regional, geo-spatial, delivery, deployment, governance & compliance, network, and premise attributes.