Network and Communication Technologies; Vol. 1, No. 2; 2012
ISSN 1927-064X E-ISSN 1927-0658
Published by Canadian Center of Science and Education

A Survey of the State of Cloud Security

Sanjay P. Ahuja[1] & Deepa Komathukattil[1]

[1] School of Computing, University of North Florida, Jacksonville, USA

Correspondence: Sanjay P. Ahuja, School of Computing, University of North Florida, Jacksonville, FL 32224, USA. E-mail: sahuja@unf.edu

Received: October 17, 2012 Accepted: November 2, 2012 Online Published: November 20, 2012

doi:10.5539/nct.v1n2p66 URL: http://dx.doi.org/10.5539/nct.v1n2p66

Abstract

Cloud computing has emerged as an important paradigm in computing today with the potential to offer scalable, fault tolerant services and reduce costs significantly. However, security concerns present significant barriers in its adoption industry wide. The multitenant nature of the cloud and the fact that data is stored in multiple locations compound these security concerns. Confidentiality, authenticity, integrity, availability and auditability are key aspects that need to be accounted for, when dealing with security. Guarantees of secure data and transactions from the service provider will enable more users to migrate to a cloud environment. Employing Intrusion Detection Systems, Cryptographic techniques and Computer Forensic Tools that recover deleted files and collect digital evidence of intruder activities are among some of the guarantees a trustful service provider can provide. This paper presents a survey on some of the common threats and associated risks on cloud platforms along with ways of tackling these threats. We also review data management and security model of some of the leading cloud service providers.

Keywords: cloud computing security, cloud computing, cloud risk assessment

1. Introduction

Cloud computing was identified by Gartner as one of the top 10 strategic technologies for 2012 with the potential for significant impact on enterprises in the coming years (Gartner, 2011). Applications deployed on the cloud inherit the ability to scale up and down giving the illusion of infinite computing resources being available. This approach allows flexibility while allocating resources and enables a cloud customer to pay only for resources that he consumes thus avoiding costs associated with over provisioning and downtime associated with under provisioning. The pay-as-you-go model is profitable to businesses that do not want to worry about maintaining the hardware or employing administration staff.

The new enterprise model of Everything-as-a-Service promotes the idea of making applications, storage and computing power available online through the cloud. Platform, Software and Infrastructure as a service are paradigms well known in computing today. Platform as a service (PaaS) provides a computing platform and environment for building web applications and services. Software as a Service (SaaS) provides business applications as a service eliminating the need for businesses to install and support applications. In the Infrastructure as a service cloud, users are given on-demand access to virtual machines. The user sees a bare bone machine with just an operating system. The user gets full flexibility to install and configure software on this machine. The same concept can be extended to Database-as-a-Service (DaaS). The past few years has seen a gradual change from in-house data management to cloud-hosted data management. Cloud Database-as-a-Service provides on-demand access to database features like data definition, storage and retrieval. It also provides data access and storage services and enables end users to be oblivious of the location and configuration of the system delivering the services.

As more and more enterprises jump onto the cloud computing bandwagon, security implications of moving services or infrastructure to the cloud need serious consideration. Small and medium businesses that typically cannot afford to devote resources to address security issues can benefit from the security solution applied by cloud providers. In order that the cloud be well accepted by organizations, security concerns of both data owners and end users need to be addressed. Moving into the cloud exposes both challenges and opportunities. Although the clouds centralized data model makes it convenient to monitor access to data, it also exposes the risk of a comprehensive data theft (Balding, 2008). Further, organizations have to trust a third party vendor with their

applications and data. This loss of control over data traditionally maintained in-house, introduces some new security management challenges. In addition, the notion of unlimited resources in the cloud is possible through resource sharing. This multitenant nature of the cloud where tenants share resources introduces new privacy concerns as the traditional network firewalls and secure socket layers cannot be a security shield in the cloud. In the cloud, a business's data is typically stored on a virtual machine, which is probably running on a server with other virtual machines some of which could potentially be malicious. In addition, cloud data is accessed via the Internet, which guarantees security only to a certain level.

2. Related Work

Security threats and risks associated with the cloud have been the focal point of many studies. Vaquero et al in (Vaquero, Rodero-Merino, & Moran, 2011) present some common cloud threats and associate these with different levels in the IaaS cloud architecture: network virtualization domain, machine virtualization domain, and physical domain. S. Subashini et al. (2011) present a survey on the security issues in different service models of the cloud: SaaS, PaaS and IaaS. They identify data security, data integrity, data availability and network security as some of the security concerns in the SaaS model. Attacks on visible code such as code running in user context are identified as an example of security issues faced in the PaaS model. Security holes in the virtualization manager are examples of issues in the IaaS model.

Minqi et al. (2010) present issues that accompany multi location of data. Many large cloud providers like Amazon and Google mirror data across geographical regions around the world in order to increase availability. This makes an already bad situation even worse because the privacy laws that apply on customers data now depend on the location of the data. The authors also review some of the privacy protection acts that came into being to protect consumers privacy but hold no ground when applied to the cloud environment. Tsai et al. (2012) discuss security threats from the perspective of virtualization technologies since virtualization is one of the foundations of cloud computing. The authors discuss the implications of virtualization on the different service models (SaaS, PaaS and IaaS). Somani et al. (2010) discuss security issues associated with storing highly sensitive medical data in the cloud. Contractual obligations committed by a service provider do not suffice when it comes to medical data management. The paper stresses on employing Computer Forensics Tools to help uncover issues with cloud.

3. Security Threats

The Verizon 2011 data breach investigations report (DBIR) presents the categories of security threats that businesses are vulnerable to today (Verizon Business, 2011). The report shows that the greatest security threats stem from external agents (92%). External agents include hackers, organized crime groups and environmental factors such as earthquakes. Insider attacks contribute to 17% of data breaches. Business partner caused breaches contribute less than 1% (Verizon Business, 2011). Business partners include suppliers, vendors, outsourced service support and any third party involved in a business relationship with the organization. The DBIR also presents the most common types of security breaches. Hacking (50%) and malware (49%) take the lead. Physical Attacks (14%) have doubled over the last couple of years. Privilege misuse (17%) like embezzlement and fraud and social tactics (11%) like solicitation and bribery fill up the bottom two positions for types of security breaches. Though the percentages shown above were not tailored specifically to cloud security threats, the report does factor in the loss of control over a company's data and assets. This loss of transparency accompanied with the shared, on-demand nature of the cloud not only introduces new challenges but also amplifies certain existing issues. In this section, we discuss some of the security challenges faced in the cloud.

3.1 Data Confidentiality

Data confidentiality deals with preventing unauthorized users from accessing data. Since enterprise data is now stored in the cloud, it is common for questions like, "Is my data safe and who else has access to it?" The cloud provider must be completely trusted to maintain security on behalf of the customer. Customers should be aware of the controls the cloud provider has on the data.

Financial and Healthcare industries require strict guarantees that their data remains protected from unauthorized access. Any breach in security can lead to penalties resulting from sensitive information disclosure and can damage the company's reputation. IaaS providers often advertise ease of registration enabling anyone to register and start using cloud services immediately. Some providers also offer free trial periods (Cloud Security Alliance, 2012). Hackers can easily target cloud providers because their lenient policies surrounding registration allows the hacker to maintain anonymity and avoid detection. IaaS and PaaS service models are the ones that are most hit by these kinds of attacks. Monitoring access and employing encryption techniques to protect data come at a significant cost in the cloud.

3.2 Insecure Interfaces

Cloud customers interact with the cloud provider through a set of interfaces or APIs that manage, orchestrate and monitor all of the customer's activity (Cloud Security Alliance, 2012). These APIs act as the gateway to cloud services. Securing this gateway has a huge impact on the overall security of the cloud services. Any unauthorized access to these gateways could compromise software integrity.

3.3 Malicious Insiders

Malicious insiders are a threat to any business organization. However, this threat is amplified in the cloud environment because a customer has no control over the hiring process of employees. The customer has no control over the kinds of background checks the cloud provider performs before hiring their employees especially the ones that have access to their data centers. Stringent checks, monitoring and logging all data access are required to help alleviate security breaches that originate from internal staff.

3.4 Shared Technology Issues

Scalability in the cloud is achieved by sharing infrastructure. Isolation between tenants is at a virtual level; the hardware and machine resources are shared. If the underlying components that make up the infrastructure are not designed for strong isolation this resource sharing introduces a potential for one tenant to peek into another tenant's data. A potential threat could arise from the type of hardware that is used in a cloud setting. Servers used in the cloud are typically multi-core and multi-processor. In such settings, information can flow between cores. In addition, cache is typically shared in multi-core processors making information stored in the cache susceptible to sniffing by a user monitoring the memory of the machine.

Virtualization is a major enabling feature of the cloud. However, it also poses security risks. Virtualization allows multiple isolated virtual machines, termed guests, to run concurrently on a single host machine. The hypervisor, also called Virtual Machine Monitor should provide perfect isolation between guests, which is lacking today. One virtualized guest environment should not interfere with another and neither should it be able to interfere with the host system. This kind of virtualization environment is susceptible to vulnerabilities like VM escape attacks and VM hopping (Owens, 2008). Virtual machines escapes allow attackers to run code on a VM that can break out and exploit a hypervisor. Once successful, the attacker gets access to the host operating system and all the other guest VMs running on that host. Exploiting the hypervisor can have serious consequences because the hypervisor runs with high privileges. Virtual Machine hopping allows one virtual machine to gain access to another compromising the victim's confidentiality, availability and integrity.

3.5 Service-Level Agreement (SLA)

A service level agreement allows a cloud customer to negotiate level(s) of service with the cloud provider. It describes the manner in which services are to be delivered. Typical SLAs include data exchange rates, mean time to repair (MTTR), jitter or similar measurable attributes of a service. Cloud customers can get security added to their SLA contracts but how do you measure security and privacy? Unfortunately, security and privacy are non-quantitative parameters making it difficult for customers to ensure that SLAs are met. Any monitoring would have to be done at the service provider end but how to ensure that the measurements provided by the service provider are accurate (Hassan, James, & Gail-Joon, 2010)?

3.6 Denial of Service Attacks

Distributed denial of service attacks exploit weaknesses in the cloud infrastructure. The cloud provider must be well equipped to be able to identify and quickly act on any high bandwidth and application layer DoS attacks. Cloud providers should have a robust infrastructure that identifies and protects against Distributed Denial of Service (DDoS) attacks. At the very least, DDoS mitigation techniques like connection limiting should be in place. DoS attacks targeted toward the shared infrastructure can enable an attacker to occupy all the available physical resources of a machine. This prevents the hypervisor from fulfilling the resource needs of other VMs (Tsai et al., 2012). The hypervisor should be configured accordingly to limit resource allocations. It is beneficial for a cloud customer to ensure that the SLAs clearly define the responsibilities of the cloud provider.

3.7 Availability

Data availability deals with prevention of malicious attacks that can make some parts or the entire application unavailable to users. Availability is not just limited to software or data, but it also includes hardware availability. Since the infrastructure is shared in the cloud, a minor glitch could result in a cloud blackout affecting even those businesses that were not the source or primary victim of the problem (PC Advisor, 2011). For example, Amazon's EC2 faced an outage in 2011 caused due to a network glitch in one of its availability zones

(Thibodeau, 2011). Many popular web sites like Netflix and Reddit hitching a ride on EC2 were affected by this outage. Cloud services should be correctly designed and maintained in order to prevent such blackouts. It is always a good idea to negotiate upfront with the cloud provider the SLAs and backup plans in the event of a cloud outage.

3.8 Compliance

Ensuring that the cloud provider follows regulations and compliance requirements as put down by requirements like HIPAA and Sarbanes-Oxley is not trivial. Requirements of HIPAA include that the patient health information should have adequate safeguards in place to ensure the confidentiality of sensitive information. Any access to patient records should be well documented and audited. Cloud providers should also be compliant with the Payment Card Industry Data Security Standard (PCI DSS). PCI DSS lays down compliance requirements to help protect cardholder information. Interestingly, 89% of victims, suffering payment card breaches, listed in the Verizon 2011 data breach investigations report (DBIR) (Verizon Business, 2011) were not compliant with PCI DSS.

3.9 Authenticity

Many companies use Lightweight Directory Access Protocol (LDAP) to manage users and groups. Active Directory is a proven and popular tool for managing users in small and medium business (SMB) companies (Subashini & Kavitha, 2011). Active Directory provides a simple authentication and security layer and allows central management of users. Companies will lose control over the management of user accounts if this functionality is hosted in the cloud, making it difficult to manage the creation and deletion of users as employees come and go. The procedure for extending existing access control mechanisms used in an enterprise up into the cloud is still not well defined.

3.10 Auditability

All accesses to a customer's instance in the cloud should be logged and audited periodically. When malware is used as the attack agent, in eight out of ten incidents, it typically sends sensitive information out of the organization (Verizon Business, 2011). It is essential to have the right indicators in place to detect this data leak. In the event of a security breach, it is necessary to have trace information that can provide evidence of foul play. The cloud environment poses an additional risk because now the cloud customer has to trust the metrics provided by the cloud provider. Another key concern is that the investigation efforts of a security breach can be hampered if the required audit logs are not made available in a timely fashion. Service level Agreements dictate how fast the cloud provider should respond in case of problems deemed urgent. The cloud provider might treat system or services being down as urgent problems, but a security breach might not enjoy the same privileges. A security breach might fall under the medium priority category. It is therefore necessary for organizations to fully understand protocols for accessing audit data when a crisis arises.

3.11 Data Integrity

Data integrity deals with preventing unauthorized modifications to data. The service provider could advertise data integrity by providing some means of encryption. Encryption can come with a performance penalty, especially when it relates to the Database-as-a-Service paradigm (DaaS). With DaaS the penalty may vary depending on if the encryption is applied to a tuple, relation, or the database as a whole. In addition, software level encryption versus hardware level encryption algorithm could make a significant impact on the performance. In software level encryption, the encryption function is provided by the database. The customer supplies the key for encryption though. This prevents unauthorized personnel from deciphering any useful information even if they get access to the disk files. On the other hand, encryption is done at the field level or row level with hardware level encryption (Hacigumus, Iyer, & Mehrotra, 2002). One problem with using this kind of fine-grained encryption is that the data size of the fields expands because of the block cipher algorithm used to encrypt data. For example, Blowfish, which is a keyed, symmetric 64-bit block cipher, encrypts and decrypts data in 64-bit chunks. If the data size of the field is 8 bytes, then after encryption by Blowfish, it will end up being 64 bytes.

With encryption in the picture, the penalty on query response time should be taken into account. The encryption query response time penalty comes from two sources: the cost involved with hardware invocation, and the encryption/decryption algorithm execution cost (Hacigumus et al., 2002). Another problem with encryption is that the key management system required to encrypt or decrypt data cannot be stored on the cloud. The customer will have to house this information. For complex encryption schemas, this key management information could take the form of a small database, which would defeat the purpose of moving the original database to the cloud.

4. Risk Assessment

Threat modeling can help identify the security threats that an application is vulnerable to along with the severity of the impact caused by the occurrence of the threat. This analysis can help identify the appropriate mitigation strategies. Security risk can be measured in terms of the probability of a threat occurring and the impact it makes. The risk is high if both the probability of the threat occurring and its impact are high (Saripalli & Walters, 2010).

Organizational risk assessment is essential during the decision making process of whether to move into the cloud or not. The level of risk involved also depends significantly on the type of cloud architecture: public or private cloud. The European Network and Information Security Agency (ENISA) divides risks into the following three categories (ENISA, 2009):

4.1 Policy and Organizational Risks

Currently there is no standardization of service interfaces or data formats provided by different cloud providers. This inherent lack of standards makes the initial push from in house data formats to the cloud provider format extremely difficult. It also makes it difficult for customers to port their applications from one cloud provider to another if the need arises. Cloud communication protocols are still very proprietary and cloud provider specific (Littlejohn, 2012). Whether cloud providers use standard protocols is important to know because a company's orchestration layer, the layer that determines how different systems used by the organization interact, should be able to support those communication protocols. This dependency on the cloud provider poses high risk if the cloud provider goes out of business or in case of acquisition of the cloud provider putting non-binding agreements at risk. Lack of standard technologies has a high probability of occurrence. The impact it has on the organization is medium (ENISA, 2009). Another instance of an organizational risk is lack of transparency in the contract. A cloud provider may outsource some of its functionalities to a third party vendor. The level of security may now depend on the level of security offered by the third party vendor. Any weaknesses in the services provided by the third party vendor can lead to adverse effect on data integrity, confidentiality and availability. If the cloud provider does not keep the customer well informed on which components of their service are outsourced, the customer might not be able to evaluate the overall risk associated with moving into the cloud.

4.2 Technical Risks

In order to project the illusion of unlimited resources in the cloud, providers typically use dynamic provisioning wherein resources like computing capacity and storage are shared among multiple users. This architecture is prone to attacks like VM hopping. The probability of this threat occurring depends on the type of cloud model; the probability is low for private clouds whereas it is medium for public clouds. The impact is high since it can affect data confidentiality, integrity and service availability. Some other high impact threats are that of a malicious insider, possibilities of interception of data in transmit and Distributed Denial of Service. The probability of these threats occurring is medium (ENISA, 2009). The multi-tenant environment also causes inconsistency for the cloud provider while determining what security controls to put in place as security requirements will vary from customer to customer.

4.3 Legal Risks

Since data is stored in multiple jurisdictions in the cloud, it would have to abide by local laws. States and countries that do not respect international agreements pose a threat to the confidentiality of sensitive data. The cloud environment makes it difficult for customers with compliance requirements to ensure that their data is in fact handled in a manner that follows those requirements. It is necessary for such customers to make sure that their cloud provider provides certification summaries of how they handle and process data. SAS70 certification is one such example (ENISA, 2009). Another legal risk involved is that the cloud provider may not report all security breaches. The probability of these threats and their impacts, both are high thus increasing the overall risk.

5. Risk Mitigation and Prevention

Appropriate security controls can mitigate risk and lower the probability of threats. It is essential for an organization to understand the risks of moving to the cloud and then have a mitigation strategy for each identified risk. One important point to remember here is that the countermeasure to a threat may not be cost effective. In this case, the risk related to that threat cannot be mitigated. Risk assessment should be done periodically because of the ever-evolving nature of the cloud. New cloud technologies can give rise to new vulnerabilities. Moreover, new countermeasures may emerge that could mitigate risks that could not be dealt with previously (Judith, 2011).

5.1 Intrusion Detection

Forensic tools and techniques can collect evidence of intruder activities that can serve as digital evidence and aid any legal procedures (Ahmed & Raja, 2010). Forensic tools can be divided into two categories i.e. tools that work on persistent data and tools that work on volatile data. Persistent tools analyze data in log files whereas volatile tools work on data that is available while the system is running.

Intrusion Detection systems (IDSs) can be used to detect attack patterns. Having such a system in place gives some level of assurance to cloud customers that the provider's infrastructure is secure. IDSs can enhance security measures by investigating user actions, network traffic, transaction logs and access logs. IDSs are a step towards proactively detecting and blocking attacks rather than dealing with the consequences of an attack after it occurs.

IDSs can be classified into host based and network based. Network IDSs monitor network traffic by capturing packets as they flow through and check them for specific patterns to identify intrusions. Host based IDSs analyze application logs, system calls and other activities on the host (Dhage et al., 2011). It can also check for activities that can only be performed via administrative access. Two different auditing techniques can be used to detect intrusions: knowledge based and behavior based (Vieira, Schulter, C. B. Westphall, & C. M. Westphall, 2010). The knowledge-based technique detects intrusions based on its knowledge of previous attacks. It looks for sequence of actions that can lead up to an attack. The behavior-based technique compares user actions to expected actions in order to detect any anomalies. Knowledge-based techniques are more common because they have a low false-alarm rate. However, only behavior-based systems can detect and capture unknown attack patterns. The ideal solution would be an IDS that uses both these techniques efficiently so as to capture a wide variety of attacks. An IDS deployed in the cloud differs from a classical enterprise version because it needs to account for internal attacks that originate from the cloud itself. In an IDS catered towards the cloud, each IDS deployed on a node monitors events for security violations and alerts the other nodes if it does encounter a violation.

5.2 Datacenter Security

In order to secure platforms, especially in IaaS, Trusted Platform Module (TPM) can be used (Vaquero et al., 2011). TPM is a specification of a secure cryptoprocessor. It is a dedicated microcontroller security chip designed to enhance software security (Parno, 2007). TPM can be used for generating cryptographic keys and has capabilities for sealed storage wherein confidential data can be bound to the platform configuration information itself. It also provides remote attestation functionality, which can be used to identify that the server is running an unmodified copy of particular software. TPM was not designed for access by multiple systems. To offset this limitation IBM developed a virtual TPM. Virtual TPM is geared toward increasing the security at all levels of the virtualization stack (Vaquero et al., 2011). TPM allows cloud providers to add several layers of authentication. TPM's security can be integrated with standards for federated ID management. Moreover, it can protect data residing in memory and disks by encrypting them. Use of TPM enabled elements by the cloud provider is assurance that the application is indeed running on trusted infrastructure.

5.3 Authentication and Identity Management

There is an increasing trend towards using multiple service providers for different aspects of the application. This can result in a potential identity management nightmare unless it is coordinated across all providers. Using federated IDM solutions is necessary to facilitate sharing of user authentication and authorization information. To ensure interoperability between communicating parties, standard protocols should be used. Security Assertion Markup Language (SAML) is an XML based open standard that allows users to log on once and exchange authentication and authorization information between different domains or different affiliated web sites (Zissis & Lekkas, 2012).

Federated solutions like User-centric IDM is a good choice for cloud environments. User-centric IDM uses attributes to identify users. The selected IDM solution should integrate with an enterprises existing IDM solution (Hassan et al., 2010). Ensure that the cloud provider has strong access controls and that the provider's security policies do not compromise what is already in house. Token or key-based authentication can be used to facilitate secure access. Digital signatures used in conjunction with Single Sign On and LDAP present a strong authentication front (Zissis & Lekkas, 2012).

5.4 Shared resources

Understanding how data is stored and accessed in the cloud will help mitigate risks associated with sharing of resources. Encrypting data before storing it in the cloud can reduce the risk of information leakage. Conventional symmetric and asymmetric encryption algorithms can be used for this purpose. In cloud environments, it is also necessary to encrypt data in cache or memory to account for sniffing attacks by neighboring nodes.

In the cloud environment, malicious parties can take advantage of the lack of isolation between tenants. In such cases, success of an attack depends on the ability of the attacker in determining co-residency of another instance (Ristenpart, Tromer, Shacham, & Savage, 2009). Inhibiting co-residency checks can thwart attacks like VM escape attacks and VM hopping attacks.

Another threat faced because of shared environment is attacks on the hypervisor. These attacks can threaten the confidentiality of tenant data. One way of protecting the hypervisor is to hide the fact that the machine is a virtual machine (Vaquero et al., 2011).

5.5 Regulatory compliance

Customers with compliance requirements should ensure that the cloud provider's design supports those requirements. For example, organizations that need to adhere to HIPAA regulations should mandate that the cloud provider support on-line probing features which provide in-depth auditing capabilities. They should be able to provide detailed activity logs that show which users accessed data along with other diagnostic information like IP addresses used to access data.

Comparing one's regulatory requirements with the cloud service providers regulatory obligations will give a good understanding of any liabilities involved with moving to the cloud. The cloud provider should be willing to undergo external audits and security certifications (Carroll, van der Merwe, & Kotze, 2011).

5.6 Location

Customers should be able to identify all possible locations of their data. Pick a service provider that can guarantee that data is stored only in geographic locations identified by the contract (Anchises, 2009). Avoid using service providers that have data centers in hostile countries. Providers should be able to prove that they are compliant with regulations and laws including location specific laws.

5.7 Availability

A good understanding of the service provider's infrastructure can identify any points of failure. Good process management can mitigate operational risk involved with cloud services. The cloud provider's data backup procedures and replication policies should be effective enough to prevent data loss or destruction (Carroll, van der Merwe & Kotze, 2011). Service level agreements that preserve uptime and provide essential workarounds in the case of failures is necessary. Service availability could be threatened in the event of the cloud provider going out of business or if the service contract with a provider ends. Countermeasure for this threat is to evaluate the provider's interoperability standards. Data extraction and data copy options should be reviewed to ensure migration to another provider will not turn out to be a nightmare.

5.8 Confidentiality

The cloud provider's security team should have adequate skills to detect and counter security breaches in a timely manner. In addition, customers should check the cloud providers breach notification process. Customers should ensure that the provider's controls around security are transparent enough and the degree of transparency should be documented in the service level agreements. Third party audits should be performed periodically to ensure policies and procedures are being followed as per specified standards.

Data in transit and data in rest can be protected using conventional cryptographic algorithms. Cloud providers should be able to attest that experienced specialists tested the encryption schemes. Key management should follow industry standards. Key loss or destruction should be prevented with adequate recovery solutions. Access to key stores should be limited to authorized personnel only (Carroll, van der Merwe, & Kotze, 2011). Periodic reviewing of privileged access is also necessary to ensure that data does not fall into wrong hands. In addition, any access to services and data should be logged and reproducible in the event of an audit.

6. Security Practices

In this section, we look at security practices and policies of some of the major cloud providers.

6.1 Amazon

Amazon offers Amazon Web Services (AWS), which is a collection of web services that together deliver a cloud computing platform. Customers can opt for an enhanced security environment with host based firewalls and host based intrusion detection systems. Employees that have potential access to customer data are subject to an extensive background check. Access rights are reviewed every 90 days and revoked if re-approval for access is not obtained. AWS maintains a service health dashboard that publishes real time service availability information (Amazon Web Services, 2012). Customers can subscribe to it via an RSS feed and get notified of any service

outages. AWS allows customers to place instances and replicate data across multiple geographic regions (Amazon, 2011). In order to account for location dependent compliance requirements, AWS lets customers pick the region associated with a geographic jurisdiction. AWS does not outsource any services to third party vendors.

AWS uses SSL-protected endpoints to deal with network security issues. To prevent packet sniffing attacks by neighboring tenants the hypervisor is hardened to send traffic to a virtual instance only if it is addressed to that instance. Amazon EC2 uses a customized version of the Xen hypervisor that is regularly evaluated for new vulnerabilities (Amazon, 2012). AWS allows customers to add additional layers of security to protect their virtual servers.

6.2 Google

Google Apps is a web-based suite of applications provided by Google as its SaaS offering. It uses a distributed file system to store data. Replication is used to distribute data over multiple systems to avoid single points of failure. Google analyses its internal traffic for presence of suspicious behavior. It also monitors system logs for unexpected activities like attempt to access customer data. Any administrative access is logged and reviewed periodically. When new employees are hired, Google conducts background checks that includes but are not limited to criminal, credit, immigration and security checks (Google, 2011). The extent of these background checks is also dependent on the job role.

Google Apps provide Single Sign-On service by exposing an API that customers can use to integrate with their in house LDAP system thus allowing companies to retain control over management of their chosen authentication mechanism. Google protects its operating systems (OS) by using proprietary software that monitors its operating systems for binary modifications (Google, 2011). Any differences found between an OS and the standard Google image of an OS triggers a self-healing mechanism that automatically restores the OS to its standard state.

6.3 Salesforce

Salesforce offers Force.com, which provides a platform for building business applications. Employees are subject to a thorough background check before being hired (Salesforce, 2012a). Secure workstations are used as a means to limit operations like cut/paste and data copying. In order to secure its networks, Force.com checks all packets flowing through the network using stateful packet inspection firewalls. In addition, Transport Layer Security (TLS) / Secure Sockets Layer (SSL) cryptographic protocols are used to encrypt all network traffic. Common external attacks are detected using Intrusion Detection Systems. Event Management tools that correlate user actions with data are used to monitor application and database activity and generate appropriate alerts (Salesforce, 2012a). Force.com maintains a trust site for real-time information on system performance and security with RSS feed capabilities (Salesforce, 2012b).

Force.com supports federated identity management using SAML, allowing affiliated web services to exchange authentication and authorization information. In addition to offering IDM solutions, Force.com also allows an organization to pick its own chosen method of authentication, for example LDAP. Network-based security is used to narrow down where users can log in from, based on the IP address they use. Force.com denies connection requests originating from unknown addresses.

6.4 IBM

IBM offers SmartCloud Enterprise, an IaaS solution that provides rapid access to virtual server environments. Provisioning of resources in the SmartCloud Enterprise environment can be done through a self-service portal or application programming interfaces (APIs) (IBM, 2011). These endpoints are subject to strict security requirements and any communication with the user is secured using SSL over HTTP. IBM offers security services for intrusion detection and vulnerability scanning. The client can use these services to monitor and scan his virtual environment. This is in accordance with IBM's shared responsibility model wherein the client is in charge of all aspects of security of his virtual environment. Clients can restrict their data to certain geographic locations by selecting the data centers where they want their resources provisioned.

7. Conclusions

There is lot of uncertainty today regarding adoption of the cloud computing model. Security concerns around data and application management are the primary contributing factors to these uncertainties. Many organizations do not want to host their data in a shared environment due to sensitivity of data, regulatory compliance or concerns with audit. Technologies provided by cloud providers are evolving to address these security concerns and make organizations more comfortable with moving their data into the cloud. From a customer perspective, it is essential to proactively evaluate all security risks of moving a business into the cloud and ensure a risk

management strategy is available for all identified risks. Cost related to a counter measure plays an important role in selecting the level of security. In this paper, we identified features a customer should look for while selecting a cloud provider. Pick a provider that can guarantee transparency of its data management activities. Draw up an SLA that clearly states the cloud provider's responsibilities in the case where security thresholds are not met.

Acknowledgement

This research has been supported by the Fidelity National Financial Distinguished Professorship in Computer and Information Sciences.

References

Ahmed, S., & Raja, M. Y. A. (2010). Tackling cloud security issues and forensics model. *High-Capacity Optical Networks and Enabling Technologies (HONET)* (pp. 190-195), 19-21.

Amazon. (2011). *Amazon Web Services Overview of Security Processes*. Retrieved from http://awsmedia.s3.amazonaws.com/pdf/AWS_Security_Whitepaper.pdf

Amazon. (2012). *AWS Risk and Compliance*. Retrieved from http://media.amazonwebservices.com/AWS_Risk_and_Compliance_Whitepaper.pdf

Amazon Web Services. (2012). *Service Health Dashboard*. Retrieved from http://status.aws.amazon.com/

Anchises, M. G. de Paula. (2009). *Cloud Computing: Enterprise Risks and Mitigation*. Retrieved from http://www.slideshare.net/anchises/cloud-computing-20091124-gts

Balding, C. (2008). *Assessing the Security Benefits of Cloud Computing*. Retrieved from http://cloudsecurity.org/blog/2008/07/21/assessing-the-security-benefits-of-cloud-computing.html

Carroll, M., van der Merwe, A., & Kotze, P. (2011). Secure cloud computing: Benefits, risks and controls. *Information Security South Africa (ISSA)* (pp.1-9), 15-17.

Cloud Security Alliance. (2012). *Top threats to Cloud Computing V1.0*. Retrieved from https://cloudsecurityalliance.org/research/top-threats/

Dhage, S. N., Meshram, B. B., Rawat, R., Padawe, S., Paingaokar, M., & Misra, A. (2011). Intrusion detection system in cloud computing environment. *In Proceedings of the International Conference & Workshop on Emerging Trends in Technology (ICWET '11)*. ACM, New York, NY, USA, 235-239.

ENISA. (2009). *Cloud Computing Risk Assessment*. Retrieved from http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment

Gartner. (2011). *Gartner Identifies the Top 10 Strategic Technologies for 2012*. Retrieved from http://www.gartner.com/it/page.jsp?id=1826214

Google. (2011). *Security Whitepaper: Google Apps Messaging and Collaboration Products*. Retrieved from http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/en/us/a/help/intl/en-GB/admins/pdf/ds_gsa_apps_whitepaper_0207.pdf

Hacigumus, H., Iyer, B., & Mehrotra, S. (2002). Providing database as a service. *Data Engineering, 2002. Proceedings*. 18th International Conference on, pp. 29-38, 07.

Hassan, T., James, B. D. J., & Gail-Joon, A. (2010). Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security and Privacy, 8*(6), 24-31. http://dx.doi.org/10.1109/MSP.2010.186

IBM. (2011). *Security and high availability in cloud computing environments*. Retrieved from http://public.dhe.ibm.com/common/ssi/ecm/en/msw03010usen/MSW03010USEN.PDF

Judith, M. M. (2011). *Cloud services: Mitigate risks, maintain availability*. Retrieved from http://www.ibm.com/developerworks/cloud/library/cl-cloudservicerisks/cl-cloudservicerisks-pdf.pdf

Littlejohn, J. (2012). *Private Cloud Blueprint.* Retrieved from http://i.techweb.com/informationweek/nwcdigital/feb12/NetworkComputing_2012_03.pdf

Minqi, Z., Rong, Z., Wei, X., Weining, Q., & Aoying, Z. (2010). Security and Privacy in Cloud Computing: A Survey. *Semantics Knowledge and Grid (SKG).* 2010 Sixth International Conference on (pp.105-112).

Owens, K. (2008). *Securing Virtual Compute Infrastructure in the Cloud*. Retrieved from http://www.savvis.com/en-us/info_center/documents/hos-whitepaper-securingvirutalcomputeinfrastructureinthecloud.pdf

Parno, B. (2007). *The Trusted Platform Module (TPM) and Sealed Storage*. Retrieved from http://www.rsa.com/rsalabs/technotes/tpm/sealedstorage.pdf

PC Advisor. (2011). *Experts explain greatest threats to cloud security*. Retrieved from http://www.pcadvisor.co.uk/news/security/3310229/experts-explain-greatest-threats-cloud-security

Ristenpart, T., Tromer, E., Shacham, H., & Savage, S. (2009). Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. *In Proceedings of the 16th ACM conference on Computer and communications security (CCS '09)*, ACM, 199-212

Salesforce. (2012a). *Secure, private, and trustworthy: Enterprise cloud computing with Force.com*. Retrieved from http://www.salesforce.com/assets/pdf/misc/WP_Forcedotcom-Security.pdf

Salesforce. (2012b). *System Status*. Retrieved from http://trust.salesforce.com/trust

Saripalli, P., & Walters, B. (2010). QUIRC: A Quantitative Impact and Risk Assessment Framework for Cloud Security. *In Proceedings of the 2010 IEEE 3rd International Conference on Cloud Computing (CLOUD '10)*, IEEE Computer Society, 280-288.

Somani, U., Lakhani, K., & Mundra, M. (2010). Implementing digital signature with RSA encryption algorithm to enhance the Data Security of cloud in Cloud Computing. *Parallel Distributed and Grid Computing (PDGC)*, 2010 1st International Conference on (pp. 211-216).

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. http://dx.doi.org/10.1016/j.jnca.2010.07.006

Thibodeau, P. (2011). *Amazon outage sparks frustration, doubts about cloud*. Retrieved from http://www.computerworld.com/s/article/9216098

Tsai, H., Siebenhaar, M., Miede, A., Huang, Y., & Steinmetz, R. (2012). Threat as a Service?: Virtualization's Impact on Cloud Security. *IT Professional, 14*(1), 32-37. http://dx.doi.org/10.1109/MITP.2011.117

Vaquero, L. M., Rodero-Merino, L., & Moran, D. (2011). Locking the sky: a survey on IaaS cloud security. *Computing, 91*, 93-118. http://dx.doi.org/10.1007/s00607-010-0140-x

Verizon Business. (2011). *2011 Data Breach Investigations Report*. Retrieved from http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2011_en_xg.pdf

Vieira, K., Schulter, A., Westphall, C.B., & Westphall, C. M. (2010). Intrusion Detection for Grid and Cloud Computing. *IT Professional , 12*(4), 38-43. http://dx.doi.org/10.1109/MITP.2009.89

Zissis, D., & Lekkas, D. (2012). Addressing cloud computing security issues. *Future Gener. Comput. Syst., 28*(3), 583-592. http://dx.doi.org/10.1016/j.future.2010.12.006