Survey of Security Issues in Cloud Computing

Nidhin K Tomson

***Abstract-* Cloud computing has quickly become one of the most prominent buzzwords in the IT world due to its revolutionary model of computing as a utility. It promises increased flexibility, scalability, and reliability, while promising decreased operational and support costs. However, many potential cloud users are reluctant to move to cloud computing on a large scale due to the unaddressed security issues present in cloud computing. In this paper, I investigate the major security issues present in cloud computing today based on a framework for security subsystems adopted from IBM. I present the solutions proposed by other researchers, and address the strengths and weaknesses of the solutions. Although considerable progress has been made, more research needs to be done to address the multi-faceted security concerns that exist within cloud computing. Security issues relating to standardization, multi-tenancy, and federation must be addressed in more depth for cloud computing to overcome its security hurdles and progress towards widespread adoption.**

Keywords--Cloud Computing; Data Security; Infrastructure; Scalability; Review

I. INTRODUCTION

Cloud Computing is a distributed architecture that centralizes server resources on a scalable platform so as to provide on demand computing resources and services. Cloud service providers (CSP’s) offer cloud platforms for their customers to use and create their web services, much like internet service providers offer costumers high speed broadband to access the internet. CSPs and ISPs (Internet Service Providers) both offer services. Cloud computing is a model that enables convenient, on-demand network access to a shared pool of configurable computing resources such as networks, servers, storage, applications that can be rapidly provisioned and released with minimal management effort or service provider’s interaction.

In general cloud providers offer three types of services i.e. Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). There are various reasons for organizations to move towards IT solutions that include cloud computing as they are just required to pay for the resources on consumption basis. In addition, organizations can easily meet the needs of rapidly changing markets to ensure that they are always on the leading edge for their consumers [1].Cloud computing appeared as a business necessity, being animated by the idea of just using the infrastructure without managing it. Although initially this idea was present only in the academic area, recently, it was transposed into industry by companies like Microsoft, Amazon, Google, Yahoo! and Salesforce.com. This makes it possible for new startups to enter the market easier, since the cost of the infrastructure is greatly diminished. This allows developers to concentrate on the business value rather on the starting budget. The clients of commercial clouds rent computing power (virtual machines) or storage space (virtual space) dynamically, according to the needs of their business. With the exploit of this technology, users can access heavy applications via lightweight portable devices such as mobile phones, PCs and PDAs.

Clouds are the new trend in the evolution of the distributed systems, the predecessor of cloud being the grid. The user does not require knowledge or expertise to control the infrastructure of clouds; it provides only abstraction. It can be utilized as a service of an Internet with high scalability, higher throughput, quality of service and high computing power. Cloud computing providers deliver common online business applications which are accessed from servers through web browser.

II. BACKGROUND OF THE STUDY

Gartner recognized seven security risks that are essential to be considered before enterprises make decisions regarding the transformation into a cloud computing model. These problems are as follows:-

1) **Authorized user access:** the potential risk of exposing organizational
data over an external processing platform, due to the limited physical, logical and personal controls outside the organizational boundaries.

2) **Conformance to regulations:** processing data outside the organizational
boundaries is still subject to accountability measures, for instance in case of auditing an external third-party space.

3) **Storage space:** cloud customer has no clue about the exact location of their data that requires service provider commitment to comply with privacy restrictions.

4) **Data separation:** clouds hold the customers’ data over a shared place where data segments are not stored in sequential manner, for that a reliable and well-tested encryption schemes are needed.

5) **Recovery:** service providers are supposed to make it clear how they will handle disasters and failures.

6) **Investigation:** Breach or intrusion attempts are hard to be tracked and spotted over the cloud due to the dispersion of the data and resources. While in some cases it could be impossible because of the high complexity level.

7) **Long-term viability:** if a rare case of service provider bankruptcy or acquisition occurs there should be a guarantee of data availability. An organization needs to be sure that it will not lose a huge amount of important data on the long-run.

Clouds as a computing model demonstrate a promising future; at the same time they highly require serious acts to cover their weak points. The weaknesses and problems come from unresolved issues in the existing technologies, which are used to build the cloud. Despite the origins or locations of risks and threats, the cloud security as an issue should be handled in a comprehensive manner [14,15]. Service providers seek fulfilling security requirements over the clouds, but face different challenges toguarantee high level of security. For that, authors in [16] discussed the requirement and challenges, also suggested

standardization and management approaches to guide cloud engineers and users. Cloud computing as an approach introduces new risks, influences others, and magnifies some. These risks and their effect on security risks and vulnerabilities were explained. Standardizing the cloud services security is an important issue that emerged due to the increased demand and importance of clouds. For instance, standardized Security Level Agreement (SLA) guarantees transparent assurance and increases the trust among cloud adopters. These standardized guarantees assist in having mutual trust, reduced risks, and better dissemination of cloud service among organizations as customers, service providers and investors.

III. CLOUD COMPUTING BUILDING BLOCKS

A. Different Service Models of cloud computing

Generally cloud services can be divided into three categories: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).

Software-as-a-Service (SaaS): SaaS can be described as a process by which Application Service Provider (ASP) provide different software applications over the Internet. This makes the customer to get rid of installing and operating the application on own computer and also eliminates the tremendous load of software maintenance; continuing operation, safeguarding and support .SaaS vendor advertently takes responsibility for deploying and managing the IT infrastructure (servers, operating system software, databases, data center space, network access, power and cooling, etc) and processes (infrastructure patches/upgrades, application patches/upgrades, backups, etc.) required to run and manage the full solution. SaaS features a complete application offered as a service on demand. Examples of SaaS includes: Salesforce.com, Google Apps.

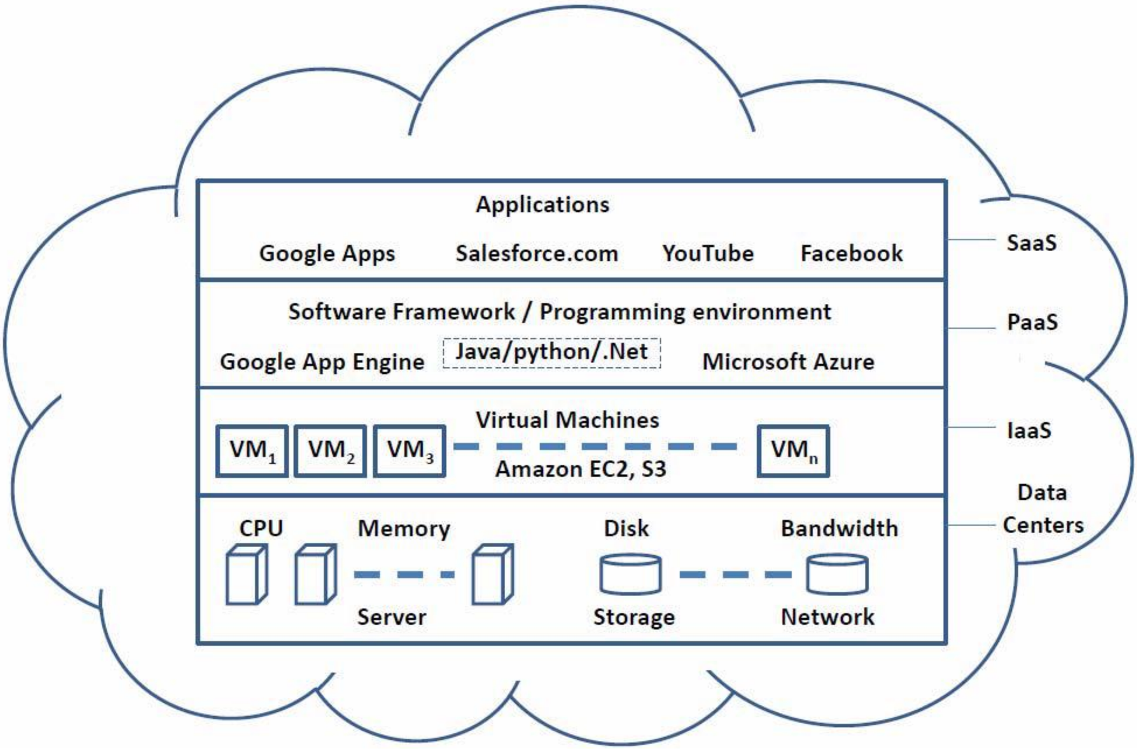

Figure 1. High Level View of Cloud Computing Architecture

Platform as a Service (PaaS): "PaaS is the delivery of a computing platform and solution stack as a service without software downloads or installation for developers, IT managers or end-users. It provides an infrastructure with a high level of integration in order to implement and test cloud applications. The user does not manage the infrastructure (including network, servers, operating systems and storage), but he controls deployed applications and, possibly, their configurations. Examples of PaaS includes: Force.com, Google App Engine and Microsoft Azure.

Infrastructure as a Service (IaaS): Infrastructure as a service (IaaS) refers to the sharing of hardware resources for executing services using Virtualization technology. Its main objective is to make resources such as servers, network and storage more readily accessible by applications and operating systems. Thus, it offers basic infrastructure on-demand services and using Application Programming Interface (API) for interactions with hosts, switches, and routers, and the capability of adding new equipment in a simple and transparent manner. In general, the user does not manage the underlying hardware in the cloud infrastructure, but he controls the operating systems, storage and deployed applications. The service provider owns the equipment and is responsible for housing, running and maintaining it. The client typically pays on a per-use basis. Examples of IaaS include Amazon Elastic Cloud Computing (EC2), Amazon S3, GoGrid.

B. *Different Service Models of cloud computing*

There are also four different cloud deployment models namely Private cloud, Public cloud, Hybrid cloud and Community cloud.

A detailed explanation about the models is given below.

Private cloud: Private cloud can be owned or leased and managed by the organization or a third party and exist at on-premises or off-premises. It is more expensive and secure when compared to public cloud. In private cloud there are no additional security regulations, legal requirements or bandwidth limitations that can be present in a public cloud environment, by using a private cloud, the cloud service providers and the clients have optimized control of the infrastructure and improved security, since user's access and the networks used are restricted. One of the best examples of a private cloud is Eucalyptus Systems.

Public Cloud: A cloud infrastructure is provided to many customers and is managed by a third party and exists beyond the company firewall. Multiple enterprises can work on the infrastructure provided, at the same time and users can dynamically provision resources. These clouds are fully hosted and managed by the cloud provider and fully responsibilities of installation, management, provisioning, and maintenance. Customers are only charged for the resources they use, so under-utilization is eliminated. Since consumers have little control over the infrastructure, processes requiring powerful security and regulatory compliance are not always a good fit for public clouds. In this model, no access restrictions can be applied and no authorization and authentication techniques can be used. Public cloud providers such as Google or Amazon offer an access control to their clients. Examples of a public cloud include Microsoft Azure, Google App Engine.

Hybrid Cloud: A composition of two or more cloud deployment models, linked in a way that data transfer takes place between them without affecting each other. These clouds would typically be created by the enterprise and management responsibilities would be split between the enterprise and the cloud provider. In this model, a company can outline the goals and needs of services. A well-constructed hybrid cloud can be useful for providing secure services such as receiving customer payments, as well as those that are secondary to the business, such as employee payroll processing. The major drawback to the hybrid cloud is the difficulty in effectively creating and governing such a solution. Services from different sources must be obtained and provisioned as if they originated from a single location, and interactions between private and public components can make the implementation even more complicated. These can be private, community or public clouds which are linked by a proprietary or standard technology that provides portability of data and applications among the composing clouds. An example of a Hybrid Cloud includes Amazon Web Services (AWS).

Community Cloud: Infrastructure shared by several organizations for a shared cause and may be managed by them or a third party service provider and rarely offered cloud model. These clouds are normally based on an agreement between related business organizations such as banking or educational organizations. A cloud environment operating according to this model may exist locally or remotely. An example of a Community Cloud includes Facebook which is showing in figure 1.

IV. FRAMEWORK FOR ANALYZING SECURITY IN THE CLOUD

Beginning in the 1980s, governmental initiatives were established around the world to define requirements for evaluating the effectiveness of security functionality built into computer systems. In 1996, initiatives from the US, Europe, and Canada were combined into a document known as the Common Criteria. The Common Criteria document was approved as a standard by the International Organization for Standardization in 1999 and has opened the way for worldwide mutual recognition of product security solutions [4].

The Common Criteria, however, serve primarily as a benchmark for security functionality in products [4]. For this reason, IBM consolidated and reclassified the criteria into five functional security subsystems. I have used these subsystems as the framework within which I assess the security issues present in cloud computing and evaluate solutions proposed.

The five functional security subsystems defined by IBM are as follows:

a. ***Audit and Compliance*****:** This subsystem addresses the data collection, analysis, and archival requirements in meeting standards of proof for an IT environment. It captures, analyzes, reports, archives, and retrieves records of events and conditions during the operation of the system.

b. ***Access Control:*** This subsystem enforces security policies by gating access to processes and services within a computing solution via identification, authentication, and authorization. In the context of cloud computing, all of these mechanisms must also be considered from the view of a federated access control system.

c. ***Flow Control:*** This subsystem enforces security policies by gating information flow and visibility and ensuring information integrity within a computing solution

d. ***Identity and Credential Management:*** This subsystem creates and manages identity and permission objects that describe access rights information across networks and among the subsystems, platforms, and processes, in a computing solution. It may be required to adhere to legal criteria for creation and maintenance of credential objects.

e. ***Solution Integrity:*** This subsystem addresses the requirement for reliable and proper operation of a computing solution.

In the next section of this paper, I address the functional systems one by one, also addressing the interactions between different functional subsystems in the section to which they most closely relate.

V. ANALYSIS OF ISSUES AND POTENTIAL SOLUTIONS WITHIN CLOUD COMPUTING SECURITY.

Audit and compliance:

Cloud computing raises issues with compliance with existing IT laws and regulations. Regulations written for IT security require that an organization using IT solutions provide certain audit functionality. However, with cloud computing, organizations use services provided by a third-party. Existing regulations do not take into account the audit responsibility of a third-party service provider.

The division of audit responsibilities required for regulatory compliance must be clearly delineated in the contracts and service-level agreements (SLAs) between an organization and the cloud provider.In order to comply with audit regulations, an organization defines security policies and implements them using an appropriate infrastructure. The policies defined by an organization may impose more stringent requirements than those imposed by regulations. It falls on the customer of the cloud services to bridge any gap between the audit functionality provided by the CSP and the audit mechanisms required for compliance.

The CSA states that the SLA between the cloud consumer and provider should include a Right to Audit clause, which addresses audit rights as required by the cloud consumer to ensure compliance with regulations and organization-specific security policies.

Even though a general approach to involve legal has been described by the CSA, no formal APIs or frameworks for integration of multiple audit systems have been defined. Additionally, there are no specific standards or models that define the separation of responsibilities between CSP and cloud service consumer.

Access control:

Access management is one of the toughest issues facing cloud computing security. One of the fundamental differences between traditional computing and cloud computing is the distributed nature of cloud computing. Within cloud computing, access management must therefore be considered from a federated sense, where an identity and access management solution is utilized across multiple cloud services and potentially multiple CSPs.

Access control can be separated into the following functions:

Authentication:

An organization can utilize cloud services across multiple CSPs, and can use these services as an extension of its internal, potentially non-cloud services. It is possible for different cloud services to use different identity and credential providers, which are likely different from the providers used by the organization for its internal applications. The credential management system used by the organization must be consolidated or integrated with those used by the cloud services.

The CSA suggests authenticating users via the consumer's existing identity provider and using federation to establish trust with the CSP. It also suggests using a user-centric authentication method, such as OpenID, to allow a single set of credentials to be used for multiple services.

Use of an existing identity provider or a user-centric authentication method reduces complexity and allows for reuse of existing systems. If done using standardized federation service, it also increases the potential for seamless authentication with multiple different types of cloud services.

The CSA states that in general, CSPs and consumers should give preference to open standards, which provide greater transparency and hence the ability to more thoroughly evaluate the security of the approach taken.

Authorization:

Requirements for user profile and access control policy vary depending on whether the cloud user is a member of an organization, such as an enterprise, or as an individual. Access control requirements include establishing trusted user profile and policy information, using it to control access within the cloud service, and doing this in an auditable way.

Once authentication is done, resources can be authorized locally within the CSP. Many of the authorization mechanisms that are used in traditional computing environments can be utilized in a cloud setting.

Federated sign-on:

A federation is a group of two or more organizations that have agreed upon standards for operation. Federations allow multiple, disparate entities to be treated in the same way. In cloud computing, federated sign-on plays a vital role in enabling organizations to authenticate their users of cloud services using their chosen identity provider.

If an organization uses multiple cloud services, it could suffer from the difficulty of having to authenticate multiple times during a single session for different cloud services. The Cloud Computing Use Cases Discussion Group suggests that the multiple sign-on problems can be solved by using a federated identity system. The federated identity system would have a trusted authority common to multiple CSPs, and provide single or reduced sign-on through the common authority.

Flow control:

Information flow control is central to interactions between the CSP and cloud consumer, since in most cases, information is exchanged over the Internet, an unsecured and uncontrollable medium. Flow control also deals with the security of data as it travels through the data lifecycle within the CSP – creation, storage, use, sharing, archiving, and destruction.

A cloud is shared by multiple service consumers, and by their very nature, cloud architectures are not static and must allow flexibility and change. Securing the flow of data across the cloud service consumer and providers and across the various components within a CSP becomes challenging and requires extensions of mechanisms used in more static environments of today.

Flow control can be separated into the following functions:-

Secure exchange of data:

Since most cloud services are accessed over the Internet, an unsecured domain, there is the utmost need to encrypt credentials while they are in transit. Even within the cloud provider's internal network, encryption and secure communication are essential, as the information passes between countless, disparate components through network domains with unknown security, and these network domains are shared with other organizations of unknown reputability.

Controls should be put in place at multiple levels of the network stack. At the application layer, I suggest using application-specific encryption techniques to ensure adequate security of the data for the particular application. At the transport layer, I suggest using standard cryptographic protocols, such as SSL and TLS. At the network layer, I suggest using network-layer controls, such as VPN tunneling, to provide easy-to-implement, secure connection with a CSP.

Data security lifecycle:

The data security lifecycle tracks the phases through which data goes from creation to destruction. It is composed of the six phases given below.

Create phase: As soon as data is created, it can be tampered with. It could be improperly classified or have access rights changed by intruders, resulting in loss of control over the data. The CSA suggests that organizations use data labeling and classification techniques, such as user tagging of data, to mitigate the improper classification of data.

Store phase: Because CSPs are third-parties, the complete security of CSP systems is unknown, so data must be protected from unauthorized access, tampering by network intruders, and leakage [10]. Due to the multi-tenant nature of cloud computing, controls must be put in place to compensate for the additional security risks inherent to the commingling of data. In order to prevent legal issues based on the physical location of data, the CSA suggests that the cloud consumer stipulate its ability to know the geographical location of its data in the SLA and ensure that the SLA include a clause requiring advance notification of situations in which storage may be seized or data may be subpoenaed.

Use and Share phase: During the use phase, which includes transmission between CSP and consumer and data processing, the confidentiality of sensitive data must be protected from mixing with network traffic with other cloud consumers. If the data is shared between multiple users or organizations, the CSP must ensure data integrity and consistency. The CSP must also protect all of its cloud service consumers from malicious activities from its other consumers.

Archive phase: As with the storage phase, data must be protected against unauthorized access by intruders, and from malicious co-tenants of the cloud infrastructure. In addition, data backup and recovery schemes must be in place to prevent data loss or premature destruction.

For data in a live production database, the CSA suggests using at-rest encryption – having the CSP encrypt the data before storage. For data that will be archived, it recommends that the cloud consumer perform the encryption locally before sending the data to the CSP to decrease the ability of a malicious CSP or co-tenant from accessing archived data.

Destroy phase: Data persistence is the biggest challenges present in the destroy phase. For data to be completely destroyed, it must be erased, rendered unrecoverable, and as appropriate, physically discarded.

The CSA suggests a plethora of techniques to be used by CSPs to ensure that data is completely destroyed, including disk wiping, physical data destruction techniques, such as degaussing, and crypto-shredding.

Identity/credentials (management):

Within cloud computing, identity and credential management entails provisioning, de-provisioning, and management of identity objects and the ability to define an identity provider that accepts a user's credentials (a user ID and password, a certificate, etc.) and returns a signed security token that identifies that user. Service providers that trust the identity provider can use that token to grant appropriate access to the user, even though the service provider has no knowledge of the user. An organization may use multiple cloud services from multiple cloud providers. Identity must be managed at all of these services, which may use different identity objects and identity management systems.

In addition, provisioning and de-provisioning of identities for an organization's IT system is traditionally done manually and infrequently. With cloud computing, access to services changes more rapidly than it would in a traditional IT application, so provisioning and de-provisioning of identities must be dynamic. Federated identity management allows an organization to rapidly manage access to multiple cloud services from a single repository. An organization can maintain a mapping of master identity objects to identities used by multiple applications within the organization's IT system. Cloud customers should modify or extend these repositories of identity data so that they encompass applications and processes in the cloud.

Currently, CSPs provide custom connectors for communication of identity and access control objects. The capabilities currently provided by CSPs are inadequate for enterprise consumers. Custom connectors unique to cloud providers increase management complexity, and are not flexible, dynamic, scalable, or extensible.

Researchers at IBM Research – China suggest using a brokered trust model, where a third-party broker server is used to establish the trust with a cloud service user. The business agreement between the CSP and the identity broker allows the CSP to place trust in the broker, allowing it to act as an agent for the CSP to establish trust with other parties, such as organizations using cloud services. The organizations can then take advantage of their own identity federation services to relay credential information for authentication with the cloud service.

Such an approach reduces the CSP's cost of establishing multiple trust relationships with multiple service users. It also pushes complexity to the trust broker, which can support more forms of federated identities. From the consumer's perspective, if multiple CSPs utilize same trust broker, establishing trust with multiple different types of services can be done by establishing trust with single trust broker.

Solution integrity:

Within the realm of cloud computing, solution integrity refers to the ability of the cloud provider to ensure the reliable and correct operation of the cloud system in support of meeting its legal obligations, e.g., SLAs, and any technical standards to which it conforms. This encompasses protecting data while it is on the cloud premises, both cryptographically and physically; preventing intrusion and attack and responding swiftly to attacks such that damage is limited; preventing faults and failures of the system and recovering from them quickly to prevent extended periods of service outage; and protection of cloud tenants from the activities of other cloud tenants, both direct and indirect.

Incident response and remediation:

Even though solutions are run by the cloud provider, cloud providers have an obligation to both their customers and to regulators in the event of a breach or other incident. In the cloud environment, the cloud consumer must have enough information and visibility into the cloud provider's system to be able to provide reports to regulators and to their own customers.

The CSA suggests that cloud customers clearly define and indicate to cloud providers what they consider serious events, and what they simply consider incidents [5]. For example, a cloud consumer may consider a data breach to be a serious incident, whereas an intrusion detection alert may just be an event that should be investigated.

Fault tolerance and failure recovery

For a CSP, one of the most devastating occurrences can be an outage of service due to a failure of the cloud system. For example, Amazon's EC2 service went down in April 2011, taking with it a multitude of other popular websites that use EC2 to host their services. Amazon Web Services suffered a huge blow from this outage. CSPs must ensure that zones of service are isolated to prevent mass outages, and have rapid failure recovery mechanisms in place to counteract outages.

The CSA recommends that cloud customers inspect cloud provider disaster recovery and business continuity plans to ensure that they are sufficient for the cloud customer's fault tolerance level.

VI. CONCLUSIONS AND FUTURE WORK

Cloud computing is an extension of existing techniques for computing systems. As such, existing security techniques can be applied within individual components of cloud computing. For example, VPN tunneling can be used for secure communication; existing encryption methods can be used to ensure protection of data on the cloud; and existing user-centric authentication methods, such as OpenID, can be used to authenticate with cloud services. However, because of the inherent features of cloud computing, such as resource pooling and multi-tenancy, rapid elasticity, broad network access, and on-demand self-service, existing security techniques are not in themselves adequate to deal with cloud security risks.

Core element of cloud computing is multi-tenancy. Due to multi-tenancy, there is a need to logically isolate the data, computing, manageability, and audit ability of users co-tenant on the same physical infrastructure at an individual component level, across architectural layers, and across multiple providers. Hence, security mechanisms and approaches that enable the above mentioned isolation in a standardized way need more scrutiny in the future.

REFERENCES

[1] National Institute of Standards and Technology, *NIST Definition of Cloud Computing*, Sept 2011.

[2] Armbrust, M. et. al., (2009), “Above the clouds: A Berkeley view of Cloud Computing”, *UC Berkeley EECS*, Feb 2010.

[3] Ramgovind, S.; Eloff, M.M.; Smith, E., "The management of security in Cloud computing," *Information Security for South Africa, 2010* , vol., no., pp.1-7, 2-4 Aug. 2010.

[4] IBM Corporation, *Enterprise Security Architecture Using IBM Tivoli Security Solutions*, Aug 2007.

[5] Cloud Security Alliance, *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1*, 2009.

[6] “Federated identity management.” Internet: http://en.wikipedia.org/wiki/Federated_identity_management, [Dec. 16, 2011].

[7] Cloud Computing Use Case Discussion Group, *Cloud Computing Use Cases Whitepaper v4.0*, July 2010.

[8] Shiping Chen; Nepal, S.; Ren Liu, "Secure Connectivity for Intra-cloud and Inter-cloud Communication," *Parallel Processing Workshops (ICPPW), 2011 40th International Conference on* , vol., no., pp.154-159, 13-16 Sept. 2011.

[9] Xiao Zhang; Hong-tao Du; Jian-quan Chen; Yi Lin; Lei-jie Zeng, "Ensure Data Security in Cloud Storage," *Network Computing and Information Security (NCIS), 2011 International Conference on* , vol.1, no., pp.284-287, 14-15 May 2011.

[10] Xiaojun Yu; Qiaoyan Wen, "A View about Cloud Data Security from Data Life Cycle," *Computational Intelligence and Software Engineering (CiSE), 2010 International Conference on* , vol., no., pp.1-4, 10-12 Dec. 2010.

[11] He Yuan Huang; Bin Wang; Xiao Xi Liu; Jing Min Xu, "Identity Federation Broker for Service Cloud," *Service Sciences (ICSS), 2010 International Conference on* , vol., no., pp.115-120, 13-14 May 2010.

[12] Shigang Chen, Meongchul Song, Sartaj Sahni, Two Techniques for Fast Computation of Constrained Shortest Paths, IEEE/ACM Transactions on Networking, vol. 16, no. 1, pp. 105-115, February 2008.

[13] King-Shan Lui, Klara Nahrstedt, Shigang Chen, Hierarchical QoS Routing in Delay-Bandwidth Sensitive Networks, in Proc. of IEEE Conference on Local Area Networks (LCN’2000), pp. 579-588, Tampa, FL, November 2000.