A Classification and Characterization of Security Threats in Cloud Computing

TARIQUL ISLAM, D. MANIVANNAN
and
SHERALI ZEADALLY
University of Kentucky, Lexington, Kentucky, USA

Security and privacy are the most critical issues that need to be addressed in designing a computing environment that is reliable and trustworthy. Like all other computing paradigms, Cloud Computing is no different. Since data and storage are outsourced to third party service providers, users lose direct control of data management and have to depend solely on the providers who may not always be dependable. This distinctive feature of Cloud Computing makes it susceptible to several security threats and vulnerabilities. Although some of the security issues such as network and virtualization security, authentication, access control, confidentiality, and integrity are not new to computing, the effect of such issues is exacerbated in cloud environment because of the unique features (e.g., multi-tenancy, data and resource sharing, virtualization, etc.) it possesses. In this paper, we classify and characterize the various security and privacy challenges associated with Cloud Computing.

Keywords: Cloud Computing, Multi-tenancy, Privacy, Security, Threat, Virtualization, Vulnerability.

1. INTRODUCTION

Due to the advantages such as flexibility and availability in obtaining computing resources at lower cost, interest in Cloud Computing has gained tremendous momentum in the last few years as observed by Armburst et al. [2009]. Cloud Computing is an abstraction based on the idea of pooling physical resources and presenting them as virtual resources. It is indeed a novel model for provisioning resources, staging applications, and platform- independent consumer access to services as mentioned by Sosinsky [2011]. One of the widely used definitions of Cloud Computing is by NIST: Mell and Grance [2011]"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction".

As observed by Hogben [2011] Cloud Computing has become one of the fastest growing paradigms of modern computing world. Since users pay only for the services they use, organizations' initial investment to adopting cloud is very low. Organizations now have the flexibility to acquire resources or services on demand. As a result, development initiatives are now at lower risk of missing the business targets as mentioned by Mather et al. [2009]. In the last few years, researchers have extensively studied various aspects of Cloud Computing including: Resource Management has been studied by Prasad and Rao [2014], Hong et al. [2015], Moreno et al. [2014], Zaman and Grosu [2013], Mastroianni et al. [2013], Papagianni et al. [2013], Morshedlou and Meybodi [2014] and Wang and Shi [2014]; Access Control has been studied by He et al. [2014], Calero et al. [2010]; Security and Auditing aspects of cloud computing has been studied by Tysowski and Hasan [2013], Xue and Hong [2014], Godfrey and Zulkernine [2014], Wang et al. [2014], Spring [2011], Wang et al. [2011] and Li et al. [2015]; issues related to Cloud Federation

Authors' addresses: Tariqul Islam and D. Manivannan, Department of Computer Science, University of Kentucky, Lexington, KY 40506. Email:{pavel.tariq@, mani@cs.}uky.edu.
Sherali Zeadally, College of Communication and Information, University of Kentucky, Lexington, KY 40506, Email: szeadally@uky.edu.

have been studied by Feng et al. [2014], Jayalath et al. [2014] and Chen and Lee [2014]. There are also several general surveys such as Heilig and Voss [2014] on cloud computing; and a survey on cloud migration by Jamshidi et al. [2013]. As observed by Chow et al. [2009], demand for cloud services is increasing at a rapid pace causing cloud service providers to overcome their limitations by creating a robust architecture to guarantee sustainable service. Quality of Service (QoS) is another important factor that needs to be met by a service provider under service level agreement as observed by CLOUD-SECURITY-ALLIANCE [2013]. Moreover, highly scalable networks, load balancing capabilities, and the ability to provide failover makes Cloud Computing services highly reliable. By outsourcing IT services to third party providers, companies can focus more on their core business Sosinsky [2011]. Despite the ever-growing interests in cloud and the plethora of services offered at a reasonable cost, Cloud Computing is susceptible to numerous threats and vulnerabilities. Several surveys such as Hogben [2011], Jansen and Grance [2011] and CLOUD-SECURITY-ALLIANCE [2013], and technical journal articles by industry experts such as Takabi et al. [2010] and Zissis and Lekkas [2012] indicate that security and privacy are the most prevailing barriers that are delaying its large-scale adoption. In Berkeley view of Cloud Computing by Armburst et al. [2009], the following ten obstacles have been identified to be hindering the widespread deployment of Cloud Computing: 1) availability of service, 2) data lock-in, 3) data confidentiality and auditability, 4) data transfer bottlenecks, 5) performance unpredictability, 6) scalable storage, 7) bugs in large distributed systems, 8) scaling quickly, 9) reputation fate sharing, and 10) software licensing. Besides these, the use of virtualization technology also introduces potential threats like hypervisor vulnerabilities, virtual machine sprawl, virtual machine side channel attacks, etc. which have been studied by Pearce et al. [2013], Ristenpart et al. [2009] and Zhang et al. [2012]. Therefore, it is crucial to have a clear understanding of the security threats associated with Cloud Computing.

Researchers are constantly identifying security and privacy loop holes in Cloud Computing. Morsy et al. [2010] have analyzed the existing challenges and grouped them according to architecture, service delivery model, cloud characteristic, and cloud stake-holder related issues. Jensen et al. [2009] have focused only on the technical security issues that arise from the usage of cloud services, especially issues related to the underlying cloud infrastructure. Subashini and V.Kavitha [2011] have presented a survey on security issues based on the service delivery models, emphasizing mainly SaaS issues. Similar work has been done by Bhadauria and Sanyal [2012] who discussed security challenges relating to the public cloud. They have also analyzed security at different levels (i.e., Network, Host, and Application level). Security and privacy challenges in Cloud Computing have been extensively surveyed by other researchers such as Grobauer et al. [2011], Gonzalez et al. [2011], Pearson and Benameur [2010] and Jansen [2011]. Although these surveys are valuable, they lack a comprehensive approach. Most of the classifications have focused on specific issues such as service delivery models, deployment models, or cloud infrastructures. Some of them did not discuss the threats associated with other distributed computing systems that can become more threatening in Cloud Computing environments. Identity management, access control, governance, legal, and compliance issues are not covered in some of the surveys. Gonzalez et al. [2011] have presented the most extensive classification on cloud security issues in recent times. They have used a quantitative approach to identify the number of references related to each category of challenges and their solutions. Thus, they have provided some insight on the issues that have received attention from the researchers and the issues which have not been talked about that much. Although they have succeeded in presenting a taxonomy of cloud security they did not delve into technical details. Therefore, a complete characterization and classification of security and privacy issues in Cloud Computing is needed. From a consumer's perspective, it is vital to identify and analyze the critical issues before deciding to outsource their sensitive data to cloud.

In this paper, we identify a variety of security and privacy threats and vulnerabilities along with some related governance and legal concerns that are considered to hinder the growth of

Cloud Computing. Following are the major contributions of this work:

—We have summarized the important issues related to security in Cloud Computing, as identified by ENISA, CSA, and NIST.

—We have identified, characterized and classified the major security threats and vulnerabilities in Cloud Computing systems. Unlike existing surveys, our classification and characterization gives a clear picture of the security threats in Cloud Computing systems which are barriers for the widespread adoption of Cloud Computing.

—We want to emphasize that we do not present a survey of existing solutions proposed for various security threats in the cloud environment. This is beyond the scope of this work.

The paper is organized as follows. In Section 2, we present an overview of Cloud Computing features that includes its key characteristics, and service and deployment models. In Section 3, we highlight three existing frameworks that focus on the critical threats and vulnerabilities in this field. In Section 4, we identify and classify the main security and privacy issues in Cloud Computing and Section 5 concludes the paper summarizing the issues presented in the paper.

2. OVERVIEW OF CLOUD COMPUTING FEATURES

According to Mell and Grance [2011], cloud model is composed of five essential characteristics, three service models, and four deployment models. In the following subsections, these features are discussed.

2.1 Essential Characteristics of a Cloud

On-demand Self-service. Computing services (e.g., server time, storage) are provisioned to meet the dynamically changing needs of the consumers.

Broad Network Access. Services are available over the network and can be accessed from heterogeneous platforms (e.g., laptops, cell phones, and PDAs) through standard interfaces.

Resource Pooling. Service providers' physical and virtual resources are dynamically allocated and de-allocated to the clients according to their changing need in a location independent manner.

Rapid Elasticity. Computing capabilities can be rapidly provisioned to quickly scale out and rapidly released as well to quickly scale in.

Measured Service. Resource usage is monitored and measured, therefore, users pay only for the services they use.

2.2 Cloud Service Delivery Models

Cloud service models are classified generally into three main categories: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). SaaS applications are intended for end-users, delivered over the web, and can be accessible from various client devices through a web browser. PaaS model provides customers tools and services to develop new applications and deploy them fast and efficiently. IaaS model provides all the hardware and software necessary to support processing, storage, networking, and other fundamental computing resources.

2.3 Cloud Deployment Models

Cloud deployment models are divided mainly into: public, private, community, and hybrid cloud. Public Cloud is available publicly and anyone can subscribe for service to the cloud. Private Cloud, on the other hand, is designed for exclusive use and accessible only within a private network. Community Cloud is owned, managed, and operated by a well-defined number of parties that have shared concerns. Finally, Hybrid Cloud is a blend of two or more cloud infrastructures (i.e. private, community, or public) bound together by standardized interfaces.

3. SECURITY ISSUES AND FRAMEWORKS

A large number of research publications addressing the security in Cloud Computing exist in the literature. Several working groups are consistently publishing articles related to Cloud Computing vulnerabilities and their mitigation techniques. Among them, the most notable ones are: CSA (Cloud Security Alliance), ENISA (European Network and Information Security Agency), and NIST (National Institute of Standards and Technology). In 2013, CSA released a report(CLOUD-SECURITY-ALLIANCE [2013]) titled The Notorious Nine: Cloud Computing Top Threats in 2013. A survey was conducted by the industry experts to identify the top threats and vulnerabilities in Cloud Computing, and eventually, nine critical threats have been listed. ENISA(Hogben [2011]) identified nine of the most important Cloud Computing specific risks in their document. Similarly, NIST, in their special publication(Jansen and Grance [2011]), highlighted the critical aspects of security. Among these issues we have identified the most critical ones and summarized them.

—Data Breaches: An organizations sensitive internal data can fall into the hands of its counterparts due to side channel timing attacks on the virtual machines. This type of attack can be designed to extract private cryptographic keys that are used in other virtual machines residing on the same physical server.

—Data Loss: Stored data can be lost due to accidental deletion, loss of encryption key, or worse, a physical catastrophe such as flood, earthquake, fire, etc.

—Account or Service Traffic Hijacking: Phishing, fraud, and exploitation of software vulnerabilities facilitate attackers to gain access to customer credentials, and that aid to launch subsequent attacks.

—Insecure Interfaces and APIs: Cloud service providers and third parties use application programming interfaces to offer different services to customers. Lack of robust identity and access management policies can lead to additional complexities and may increase the risk as well.

—Denial of Service: Attackers generate huge amount of fake requests to a certain cloud server so that the server is forced to consume processor power, memory, disk space, and network bandwidth. This eventually causes an intolerable system slowdown and keeps other customers off the service.

—Malicious Insiders: System administrators, current or former employees, contractors, or other third party service providers who have or had the access privilege may misuse this and may cause intentional damage, affecting the confidentiality, integrity, and availability of an organizations sensitive data.

—Insufficient Due Diligence: Lack of proper understanding of the service provider environment and improper assessment of the operational ins-and-outs may put an organization into critical situation if it rushes to adopt the technology.

—Shared Technology Vulnerabilities: Hypervisor vulnerabilities, cross VM side channel attacks, VM sprawl, and many other possible threats due to the shared multi-tenant architecture can expose the entire environment to a potential compromise.

—Loss of Governance: Transferring data to the cloud means transferring the control to the service provider. On a number of issues this may have security implications.

—Lock-in: Since there is no well-established standard for data and service portability, dependency on a particular cloud service provider often makes it tough for the client to migrate from one provider to another.

—Insecure or Incomplete Data Deletion: Deleting data from the cloud storage does not guarantee that it will be inaccessible in future. In fact, if data is not securely erased by the provider or the disk is not encrypted by the client, data could be reconstructed later.

—Availability Chain: A cloud service provider can delegate some of its work to a third party or even can use the service of another service provider. Thus, a potential for cascading failures is

created that may affect service availability.

Next, we classify and characterize the key security threats and vulnerabilities in Cloud Computing.

4. KEY SECURITY THREATS AND VULNERABILITIES

As mentioned previously, security is considered to be one of the most important barriers for widespread adoption of Cloud Computing. Therefore, understanding the security issues in Cloud Computing and devising efficient solutions are critical for its success. In this section, we identify the main security threats and vulnerabilities associated with Cloud Computing and group them into six categories: Network Security, Virtualization and Hypervisor Security, Identity and Access Management, Data and Storage Security, Governance, and Legal and Compliance issues.

4.1 Network Security

Issues related to security in Cloud Computing include issues that exist in traditional computing environment and issues specific to Cloud Computing. In this section, we present a selected list of security issues associated with network communications and configurations in Cloud Computing environments.

4.1.1 *XML Signature Wrapping Attack.* XML signatures are widely used for the purpose of authentication and integrity of SOAP (Simple Object Access Protocol) messages. Protocols that use XML signatures suffer from a well known attack called XML Signature Wrapping Attack or simply Wrapping Attack ad mentioned by Jensen et al. [2009]. This applies to Web Services and therefore also for Cloud Computing. To ensure message integrity, a predefined part or parts of the SOAP message are signed using XML signature. The message contains a security header with a signature element, which references one or more parts of the message that have been signed. An XML Signature Wrapping Attack essentially exploits the fact that the signature element does not convey any information about the referenced part of the message. An attacker can easily modify the message body and inject malicious code without invalidating the signature. Here, the attacker virtually wraps the XML signature around the malicious code and passes it on as if it were an authentic message(II and Al-Hamdani [2011] and Kouchaksaraei and Chefranov [2013]).

4.1.2 *Flooding Attacks.* Whenever a company's computational demand increases, more instances of virtual machines are assigned instantaneously by the cloud to cope up with the demand. But this also opens the door for malicious adversaries to exploit this feature as observed in Jensen et al. [2009]. By instantiating a large number of virtual machines, an attacker can generate a huge amount of fake requests and forward them toward a certain server. But in order to determine the validity, the server must inspect each and every one of those requests. As a result, the entire network gets flooded with requests, and legitimate requests starve; it leads to a distributed denial of service (DDoS) attack(Zunnurhain et al. [2014]).

4.1.3 *Malware Injection Attack.* This type of attack involves injecting a malicious service implementation or installations of virtual machine instances into the cloud system(Jensen et al. [2009]). In SaaS or Paas environment, the attacker's goal is to create his/her own service implementation containing malicious code or scripts and to deploy that service in the cloud so that it looks like a legitimate service. If he/she is successful, this kind of Malware could serve any malicious operation they intend to perform. Thus, when valid requests arrive, the cloud system automatically redirects them to such malicious service implementation. The impact of this attack includes eavesdropping, subtle data modification, unauthorized access to cloud resources, user credential leakage, functionality changes, and service blocking as described by Jensen et al. [2009] and Zunnurhain and Vrbsky [2010]. In the same way, in the IaaS environment, attackers can instantiate virtual machines and inject Malware into them. Legitimate users' requests starve until the fake services are completed. This can lead to service deadlock if the number of requests

is huge(Zunnurhain and Vrbsky [2010]). The key challenge here is not just detecting the Malware Injection Attack but also determining the virtual machine instances that are used by the attacker for the malicious service implementation(Khalil et al. [2014]).

4.1.4 *Metadata Spoofing Attack.* When communicating with other Web services, a Web service client needs to retrieve all necessary information regarding a Web Service invocation. This information includes Web service address, message format, network location, security requirements, etc., which are stored in the meta data documents provided by the Web services server. Two of the most common Metadata documents are Web Service Definition Language (WSDL) file and WS-Security-Policy. Jensen et al. [2009] observe that Metadata documents are likely to open the possibilities of spoofing attacks since they are distributed using communication protocols like HTTP or e-mail they. It is possible for attackers to maliciously alter the content of the WSDL file and distribute them across all the Web service clients. This has serious consequences and security implications. As explained by Grobauer et al. Grobauer et al. [2011], an adversary may modify a WSDL file in such a way that a call to a *deleteUser* operation syntactically looks like a *setAdminRights* operation. When a user is provided with this altered WSDL file, each of his/her *deleteUser* operation invocations will be replaced by the *setAdminRights* operation invocations. Thus, users who are supposed to be deleted now become active with administrator privilege. Another way WSDL spoofing can be done is by modifying the network endpoints and the references to security policies (WS-Security-Policy). Modified network endpoints facilitate an adversary to easily launch a man-in-the-middle-attack(Jensen et al. [2009]). Metadata spoofing can be dangerous in Cloud Computing environment, where the cloud system itself has some WSDL repository functionality. It is assumed that new users can dynamically retrieve a service's WSDL file and spread the malicious WSDL file throughout the network(Grobauer et al. [2011]).

4.1.5 *Insecure APIs.* Cloud customers usually use a set of Application Programming Interfaces (APIs) to manage and interact with the cloud services. Service provisioning, application management, monitoring, enabling security functions, etc., are all performed through these interfaces(CLOUD-SECURITY-ALLIANCE [2013]). These APIs are critical for the security and availability of the cloud service. If they are not designed with robust identity and access management policies they could be easy target for the malicious attackers. Attackers always try to exploit the security vulnerabilities in these APIs by circumventing the policies(Grobauer et al. [2011]). Moreover, organizations and third parties often use these APIs to offer different value-added services to their customers. This new layered API makes things more complicated because organizations may need to share their credentials to third party service providers to enable those services. This has serious security and privacy implications from the organizations' perspective. Although cloud service providers always attempt to ensure that security is well integrated into their service models, consumers need to understand and analyze the security implications before outsourcing any service. Unsecured and poorly designed APIs can expose an organization to a variety of security threats that will impact confidentiality, integrity, availability and accountability(CLOUD-SECURITY-ALLIANCE [2013]).

4.1.6 *Cross Site Scripting (XSS) Attack.* Cross Site Scripting (XSS) attacks basically inject malicious scripts or codes into Web contents and thereby force a website to execute the attacker's supplied codes. The end user is the intended victim and the attacker exploiting the vulnerable website acts as the media for this type of attacks(Bhadauria and Sanyal [2012] and Hydara et al. [2015]). Failure to properly validate user input is the root cause of XSS attacks. In this case, two things can happen, either the web site fails to neutralize the user input or it does the validation incorrectly. Thus, this opens the door for the attackers to exploit the vulnerabilities. Attackers can steal Web browser cookies and hijack online user credentials, extract sensitive user data, and also perform many other malicious activities. It has been discovered in recent research that XSS attackers can gain full control over the Web browser, similar to Trojan-horse programs(Hydara et al. [2015]). Users can be affected by XSS in two ways. While browsing

the Internet some specially crafted link or popups can open up on the screen and users can be either tricked into clicking that link, or they can be attacked while merely visiting a Web page embedded with malicious code in it. Thus, the intruding party gets control over the user's private data Bhadauria and Sanyal [2012].

4.1.7 *SQL injection Attack.* SQL injection attacks are the class of attacks in which malicious code is inserted into the data fields of a standard SQL query. Thus attackers gain unauthorized access to databases(Bhadauria and Sanyal [2012]). A successful exploit allows attackers to extract private and sensitive data from the database, tamper existing data by modifying the database through insert, delete, or update operations. It is even possible to modify the roles/privileges of the users and execute administration operations that can lead to complete destruction of data form the database server. Use of dynamically generated SQL query and inadequacies in handling user input are the key reasons for such attacks(Liu et al. [2009]).

4.2 Virtualization and Hypervisor Security

Virtualization is one of the core components of Cloud Computing that helps organizations optimize their application performance in a cost effective manner(Sabahi [2012]). This technology can be used as a security component also; for instance to provide monitoring of virtual machines, facilitating management tasks such as performance management, cloud infrastructure management and capacity planning management(Lombardi and Pietro [2011]). Hypervisor acts as the abstraction layer providing necessary resource management functions to enable sharing of hardware resources between the virtual machines(Pearce et al. [2013]). Although there are great benefits to be gained from these technologies, they also introduce additional security threats which are discussed next.

4.2.1 *Hypervisor Vulnerabilities.* A hypervisor or Virtual Machine Monitor (VMM) is designed to run multiple guest VMs and applications concurrently on a single host machine and to provide isolation between the guest VMs. Although hypervisors are expected to be robust and secure, they are vulnerable to attacks. If attackers get control of the hypervisor, all the VMs and the data accessed by them will be under their full control to exploit(Morsy et al. [2010]). Another reason hackers consider the hypervisor a potential target, is the greater control supported by the lower layers in the virtualized environment. Compromising a hypervisor also helps to gain control of the underlying physical system and the hosted applications. Some of the well-known attacks (e.g., BLUEPILL, Hyperjacking, etc.) insert VM-Based Rootkits that can install rogue hypervisor or modify the existing one to take complete control of the environment(Ibrahim et al. [2010]). Since hypervisor runs underneath the host OS, it is difficult to detect these sorts of attack using regular security measures.

4.2.2 *VM Escape.* Virtual machines are designed to support strong isolation between the host and the VMs. But the vulnerabilities in the operating system running inside the VMs can aid attackers to insert malicious program into it. When that program is run, VM breaks the isolated boundaries and starts communicating with the operating system directly bypassing the Virtual Machine Monitor (VMM) layer. Such an exploit opens the door to attackers to gain access to the host machine and launch further attacks (Luo et al. [2011]).

4.2.3 *VM Sprawl.* VM sprawling occurs when a large number of virtual machines exist in the environment without proper management or control. Since they retain the system resources (i.e., memory, disks, network channels etc.) during this period, these resources cannot be assigned to other VMs, and they are effectively lost(Luo et al. [2011]). Dabrowski and Mills [2011] demonstrate two circumstances that can cause VM sprawling and contribute to the creation of orphan VMs. VMs are usually allocated and terminated upon request from the users, and the system generates the acknowledgement messages in response. The problem arises when VM creation or termination is completed but the messages are lost in transit. Users retry by generating new requests until they become successful, and this causes orphan VMs to grow in number.

Eventually system resources get exhausted and that leads to a collapse in the overall performance of the system. Migrating the orphan VMs to another lightly-loaded physical server may solve the issue to a certain extent. But, ensuring the same level of security configurations, Quality of Service (QoS), and enforcing privacy policies is always a challenge as observed by Sabahi [2012].

4.2.4 *Cross VM Side Channel Attack.* To maximize resource utilization, multiple VMs are usually placed on the same physical server in the cloud environment and this co-resident placement is a potential threat to cross VM side channel attack. The basic idea is: a malicious VM penetrates the isolation between VMs, and then access the shared hardware and cache locations to extract confidential information from the target VM. Ristenpart et al. [2009] first showed that it is possible to map the internal cloud infrastructure and deliberately place a malicious VM onto the same physical server the target VM is likely to reside. Having placed the malicious VM co-resident with the victim VM, they showed preliminary results on a variety of cross-VM side channel attacks, including denial of service (DoS), remote keystroke monitoring via timing inference and others. Zhang et al. [2012] demonstrated that, by launching an access-driven side-channel attack, it is possible for a malicious virtual machine to extract fine-grained information (i.e., private key) from a victim VM running on the same physical server. Similar results have been reported by Apecechea et al. [2014] and they showed that fine-grained cross-VM attacks are possible in modern virtualized servers.

4.2.5 *Outdated SW Packages in VMs.* Outdated software packages in virtual machines can pose serious security threats in virtualized environment. Because of the low cost and the ease of creation, users tend to create new virtual machines for different tasks, branch new virtual machines based on the old ones, snapshot machines or even rollback machines to an earlier state. These operations may have serious security implications, for example, a machine rollback operation may expose a software bug that has already been fixed as mentioned in Schwarzkopf et al. [2009]. In Infrastructure-as-a-Service cloud, service providers usually delegate the task of updating the software packages to the clients. But this can create problems if clients fail to do that, resulting in a large number of outdated virtual machines that are not patched for the recently discovered software vulnerabilities(Schwarzkopf et al. [2011]). The issue becomes more severe if the number of VMs running in a virtualized infrastructure increases over time. User gets overwhelmed to keep the software packages up to date. Furthermore, virtual machines that remain inactive for long period of time are difficult to be updated with the latest security patches since this would require the machines to be restarted from the dormant state first(Schwarzkopf et al. [2012]).

4.2.6 *Single Point of Failure.* Existing virtualized environments are based on the hypervisor technology that controls the access of the VMs to physical resources and is important for the overall functioning of the system. Therefore, failure of the hypervisor due to overused infrastructure or software faults leads to the collapse of the overall system as observed by Sabahi [2012]. The entire system must be rebooted to recover from such failures. But, such failures leads to arbitrary state corruptions and inconsistencies throughout the system, and hence all the works that were in progress in all the VMs are lost(Le and Tamir [2011]). Another drawback of the virtualized servers is that they have a finite number of access points (i.e., Network Interface Cards) for all the VMs. Compromising these access points could open the gateway for the attackers to exploit the virtual cloud infrastructure including the VMs, hypervisor, and the host machine as mentioned by Ibrahim et al. [2010].

4.2.7 *VM Image Sprawl.* Secure management of VM images is an important requirement in Cloud Computing environment. Each VM image is actually a full software stack containing installed and configured applications that are used to boot the VM into an initial state or the state of some previous checkpoint (Jansen [2011]). They are treated as data, and therefore easy to clone, extend, and snapshot. Thus, VM images grow in number and can cause VM image sprawl

(Reimer et al. [2008]). Another critical issue is that many of the VM images are designed to be shared by different and often unrelated users. If image repositories are not carefully managed and controlled, sharing of VM images may pose privacy and security threats as observed by Wei et al. [2009]. Moreover, since an image can contain propriety code and data, the owner of the image risks releasing sensitive information inadvertently. Attackers would be keen on examining the image to discover security loopholes. Attackers may also inject malicious codes in the VM image, or supply a completely new VM image containing a Malware as mentioned by Jansen [2011] and Morsy et al. [2010].

4.3 Identity and Access Management (IAM)

Managing identities and providing secure and efficient access to large-scale outsourced data is an important element of Cloud Computing (Wang et al. [2009] Yu et al. [2010]), and this remains one of the greatest challenges the IT industry is facing today (csa []). For most organizations, data security and privacy issues are most important. So, a good identity and access management strategy is a necessary prerequisite for strategic use of secure on-demand Cloud Computing services. The proofs of user identity and authentication aspects of identity management involve the use, maintenance, and protection of personally identifiable information (PII) collected from cloud consumers. Therefore, Jansen and Grance [2011] observe that thwarting unauthorized access to data resources in the cloud also deserves a major attention.

Following are the major IAM challenges that need to be addressed for successful and effective management of identities in the cloud.

Identity Management:. Secure and efficient management of provisioning and deprovisioning of users to systems and applications is one of the major challenges in Cloud Computing (csa []). Mather et al. [2009] observe that frequent changes in users' roles and responsibilities inside the organization, turnover of users, changes in business (e.g., mergers and acquisitions, process outsourcing) are the factors that affect establishing a sustainable IAM process.

Authentication:. Authenticating the identity of a user or a system in a secure and dependable way is another key issue. Other challenges include proper credential management, ensuring robust authentication, compliance with the password standard, encryption management, and managing trust across all types of cloud services (Mather et al. [2009] and csa []).

Authorization and Access Control:. Establishing fine-grained authorization and access control policies for users to access the systems resources (i.e., applications, databases, etc.) is another vital requirement csa []. Jansen and Grance [2011] observe that adapting to the continuous changes in users' roles or privileges and maintaining control over access to resources are also challenging.

Federation Management:. Federated identity management lets organizations authenticate their users (providing single sign on facility) by exchanging identity information between the the Service Provider (SP) and the Identity Provider (IdP) (Yan et al. [2009]). Since identity information are dynamically distributed across security domains, it poses significant security and privacy challenges. Insecure communication network and weak user authentication scheme in Web identity chain can lead to replay attacks, session hijacking, and phishing attacks as observed by Maler and Reed [2008]. Furthermore, Hackett and Hawkey [2012] observe that reliance on IdP for identity management may cause identity theft and data breaches if the IdP behaves maliciously.

4.4 Data and Storage Security

Ateniese et al. [2008] observe that the concept of third-party data warehousing, more commonly, data outsourcing has become a rising trend. Therefore, users are to depend solely on the respective cloud providers for the availability, integrity and confidentiality of their data (Wang et al. [2009]). A cloud service provider that stores consumers data is not necessarily trusted. Therefore, the issue of ensuring the integrity of the stored data at untrusted servers has received a lot of attention as mentioned by Erway et al. [2009] and Zeng [2008]. Moreover, it is quite possible that storage

service providers may decide to hide the data loss incidents from the customers to keep their reputation undamaged. To aggravate the situation even more, providers may sometimes overlook the importance of sensitive customer data and end up deliberately deleting rarely accessed files, especially those from ordinary clients. So, despite being envisaged as a very promising service platform for modern computing paradigm, this novel data-storage model in cloud brings forward many thought-provoking topics that have great influence on the security and the performance of the overall system (Wang et al. [2009] and Wang et al. [2011]). That is, even though third-party data outsourcing into the cloud is economically beneficial, lack of strong assurance of data confidentiality, integrity and availability is hindering the large-scale adoption by organizations as well as end users.

Security issues that can arise due to data and storage outsourcing are discussed next.

4.4.1 *Data Confidentiality.* Confidentiality refers to limiting information access and disclosure only to authorized users or system. One of the fundamental principles of confidentiality is "need-to-know" or "least privilege" (Zissis and Lekkas [2012]). In effect, access to vital and sensitive information should be restricted only to those individuals or systems that have a specific need to get or use that information. In Cloud Computing environment, due to the large number of parties, devices and applications involved, the number of access points also increases. Therefore, the risk of data breaches increases as well. The potential concerns that might affect confidentiality of the data stored in a public cloud are: 1) access control (authentication and authorization) mechanisms, 2) data protection scheme, 3) encryption algorithm used, and 4) encryption key management.

4.4.2 *Data Integrity.* Data Integrity in Cloud Computing is considered one of the biggest concerns. Integrity means information is accurate and reliable and has not been subtly altered or tampered by an unauthorized party. The term integrity in fact is associated with authenticity: the ability to verify that content has not been changed in an unlawful manner; and non-repudiation and accountability: the source of any action performed on the system can be verified and associated with a user. In addition to ensuring the confidentiality of the data, cloud consumers also need to worry about the integrity of their data. In the case of confidentiality, use of any strong encryption method alone is sufficient while integrity requires the use of message authentication codes (MACs) (Mather et al. [2009]).

4.4.3 *Data Availability.* It is expected that data and other critical assets would be accessible to customers and businesses when needed. Despite employing the architectures designed for high service reliability and availability, Cloud Computing services can and do experience outages and performance slowdown. Availability can be affected temporarily or permanently, and data loss can be partial or complete. Quite a number of threats that can hamper the availability exist. Firstly, network-based attacks such as the denial of service (DoS) attack can affect availability. Secondly, cloud service providers own availability can be another important concern. In addition to service outages, largescale storage systems can experience disk/sector failures, some of which can result in permanent data loss. And with the exponential growth of archival data, a small failure rate can imply significant data loss in archival storage (Chen and Lee [2014]).

4.4.4 *Data Isolation.* Multi-tenancy and shared resources are the vital characteristics of Cloud Computing. Due to the presence of multi-tenants in a cloud environment, resources (i.e., servers, storage) are shared by multiple organizations that provide flexibility and economies of scale. But, from a customers standpoint, the notion of using shared infrastructure could be of great concern. The concentration of data and resources introduces various associated risks, including sharing the infrastructure with untrusted tenants and relying on the availability and security of the underlying infrastructure itself. Mundada et al. [2011] observe that these security vulnerabilities represent some of the most significant obstacles to the adoption of cloud-based services. Therefore, before moving their data to cloud, administrators need to ensure that all data in cloud are completely

secure and accessible only by authorized users. Usually, in a cloud environment, a customers request is processed by an application that runs with adequate privileges to access any tenants data any time. This application is responsible for authenticating and authorizing request. As the only protection is at the application level, a single vulnerability at this level threatens the data of all tenants which could also lead to cross-tenant data leakage, making the cloud much less secure than dedicated physical resources as mentioned by Factor et al. [2013].

4.4.5 *Data Sharing.* Due to its intrinsic data and resource-sharing nature, cloud-based services are an attractive model for user-facing applications like online word processing, calendaring, blogging, and social networking. These applications allow multiple users to edit their shared resources concurrently, while being scalable, highly available and globally accessible (Feldman et al. [2010]). These benefits on the contrary can affect privacy due to server-side information leakage and pose significant risk to the confidentiality of those shared resources (Chu et al. [2014]). Storing data in the cloud alone is not adequate; it might also be essential to guarantee anonymity. But, unrestricted anonymity can cause serious problems also. A defied employee of an organization can mislead others by sharing false files without even being traceable as observed by Ruj et al. [2014]. Lastly, the recurrent change of membership in a group makes it difficult to share data in a multi-owner environment while preserving data integrity and privacy at the same time as mentioned by Liu et al. [2013].

4.4.6 *Data Backup & Redundancy.* Outsourcing data to the cloud storage does not necessarily mean that data is actually backed up. Data could be lost accidentally, modified by adversaries, and even encryption key could be lost. If the original copy of the data is not properly backed up, recovery would be impossible. To avoid data loss and maintain business continuity consumers must ensure that proper backup policies are in place. Because of the ease of operations, service providers may prefer to rely on seamless backups without the active consent of the clients (Pearson and Benameur [2010]). Tang et al. [2010] observe this approach is undesirable since data can be inadvertently disclosed in future due to some external or internal attacks on the cloud or erroneous management by the cloud operators. Rahumed et al. [2011] observe that controlling the version of the backups is another challenging issue since data could be replicated multiple times by the service provider over the infrastructure.

4.4.7 *Data Sanitization.* Data sanitization is the process of expunging data from the storage media so that data cannot be reconstructed later. In public cloud environment, complete deletion of data (upon request from the client) including all the log files and backup replicas made for recovery is a fundamental obligation as observed by Gonzalez et al. [2011]. But, timely destruction of data might be challenging since multiple replicas of the data could be dispersed in different geographical locations. Therefore, it is difficult to ensure whether a service provider is reliably removing all backup copies of the data. Moreover, the disk that needs to be destroyed may also share data from other clients. Sometimes destroying the storage media itself can be a necessity to ensure complete deletion of data. If they cannot be disposed properly, it might be possible to reconstruct data from those abandoned media.

4.4.8 *Data Provenance.* Data provenance security in Cloud Computing is an area of significant concern that deserves careful attention. Provenance basically refers how data has been generated, who has accessed and modified the data, and what the sequences of those actions are. The provenance of sensitive data may divulge critical private information, and adversaries always look for security loopholes to exploit this. Data provenance could be valuable in cases where information trace-back, auditing, forensic analysis, and history-based access control is needed. Revealing data provenance on the contrary can impact data privacy and maintaining a balance between these two is considered a major challenge. Hence, in addition to protecting integrity of the sensitive data, it is essential to make the data provenance secure (Takabi et al. [2010] and Asghar et al. [2011]).

Table 1: A Classification of the Security Threats and Vulnerabilities in Cloud Computing

No.	Category	Security threats and Vulnerability
1	Security at Network Level	XML Signature Wrapping Attack Flooding Attack (DDoS) Malware Injection Attack Metadata Spoofing Attack Insecure Web Applications and APIs Cross Site Scripting Attack SQL Injection Attack
2	Virtualization Security	Hypervisor Vulnerabilities VM Escape VM Sprawl Cross VM Side Channel Attack Outdated SW Packages in VMs Single Point of Failure VM Image Sprawl
3	Identity and Access Management	Identity Management Authentication Authorization & Access Control Federation
4	Data and Storage Security	Data Confidentiality Data Integrity Data Availability Data Isolation Data Sharing Data Backup and Redundancy Data Sanitization Data Provenance Dynamic Data Updates
5	Governance	Improper Data Sanitization Information leakage Vendor Lockin
6	Legal and Compliance Issues	Data Location Contracts and Electronic Discovery Laws and Regulations Audit Assurance

4.4.9 *Dynamic Data Updates.* Ensuring remote data integrity in Cloud Computing environment is a difficult task, especially when data is frequently updated by the clients through block modification, insertion, and deletion. Most of the existing works on remote integrity checking focus on static archive data and therefore cannot be applicable to cases where dynamic data updates are more common. Furthermore, direct extension of the current Provable Data Possession (PDP) or Proof of Retrievability (PoR) techniques that can support data dynamics may lead to security loopholes. Therefore, an efficient and provably secure dynamic auditing protocol is highly desirable in cloud environment to verify the integrity of the data(Yang and Jia [2013] and Wang et al. [2009]).

4.5 Governance

In cloud environment, consumer relinquishes control to the cloud service provider on a number of critical issues (e.g., policies, procedures, and security mechanisms of deployed services) that have security implications. Following are the issues that stem from this loss of governance:

Improper Data Sanitization:. If data is not securely erased by the service provider, data could be reconstructed later from the disk (considering disks are not encrypted by the client).

Data and Information leakage:. From consumers perspective, transferring data to the cloud means giving up control over the data backup procedures, file systems, redundancy, security policies, and other relevant configurations.

Vendor Lock-in:. No firmly-established standard exists for data and service portability in Cloud Computing environment yet. Therefore, if a consumer becomes dependent on a particular service provider then it would be difficult to migrate to another service provider.

4.6 Legal and Compliance Issues

Legal and compliance aspects refer to the responsibilities of an organization that are essential to operate in accordance with established laws, regulations, standards, and specifications. Jansen and Grance [2011] and sec [2009] have have identified the following as potential concerns in this area:

Data Location:. Data can be stored redundantly in multiple geographical locations and detailed information about the data location may not be disclosed to the client. That means, when data crosses borders, the governing, legal, compliance, and regulatory administrations can be ambiguous and raise a variety of other security concerns.

Contracts and Electronic Discovery:. Legal issues may arise when dealing with electronic discovery that involves the identification, collection, and analysis of stored data in the discovery phase of a litigation.

Laws and Regulations:. Different countries have different types of security and privacy laws and regulations at various levels (i.e., local, national, state, etc.) which makes legal and compliance issues more complicated.

Audit Assurance:. It is important to ensure audit assurance for a proper organizational governance. Designing a robust audit methodology to reflect the various compliance requirements is always a tricky task.

A summary of the security threats and vulnerabilities discussed in Section 4 is presented in Table 1.

5. CONCLUSION

In this paper, we discussed the essential characteristics of cloud, its service delivery and deployment models, compelling reasons for adopting it, and the barriers that hinder its wide adoption. We also surveyed three well-known cloud security frameworks namely ENISA, CSA, and NIST that aim to provide a compilation of risks, vulnerabilities and also the best practices to resolve them. These three entities provide a comprehensive overview of the current security, privacy, and trust issues, and thus, help in understanding the current status of cloud security. Then, we presented a variety of security and privacy concerns associated with Cloud Computing, identified major threats and vulnerabilities, and classified them into six categories: Network Security, Virtualization and Hypervisor Security, Identity and Access Management, Data and Storage Security, Governance, and Legal and Compliance issues. Each of these categories identified several threats and vulnerabilities, resulting in further classification. It is evident from our discussion that for the wide spread adoption of the cloud, these issues must be addressed thoroughly. Therefore, enrichment of the existing solution techniques as well as more innovative approaches to mitigate these problems are needed. Though Cloud Computing is a hot area, it is still in its infancy, and its widespread adoption will depend mostly on how the ever increasing security concerns will be addressed in the upcoming days.

REFERENCES

CSA DOMAIN 12. https://cloudsecurityalliance.org/guidance/csaguide-dom12-v2.10.pdf.

2009. Security Guidance for Critical Areas of Focus in Cloud Security Computing V3.0. http://www.cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf.

Apecechea, G. I., Inci, M. S., Eisenbarth, T., and Sunar, B. 2014. Fine Grain Cross-VM Attacks on Xen and VMware. In *Proc. BDCloud*. 737–744.

Armburst, M., Fox, A., Griffith, R., Joseph, A., Katz, R., Konwinski, A., Lee, G., Peterson, D., Rabkin, A., Stoica, I., and Zaharia, M. 2009. Above the Clouds: A Berkely View of Cloud Computing. Tech. Rep. UCB/EECS-2009-28, University of California at Berkely, eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html.

Asghar, M. R., Ion, M., Russello, G., and Crispo, B. 2011. Securing Data Provenance in the Cloud. In *Proc. iNetSeC*. 145–160.

Ateniese, G., Pietro, R. D., Mancini, L., and Tsudik, G. 2008. Scalable and Efficient Provable Data Possession. In *Proceedings of SecureComm*.

Bhadauria, R. and Sanyal, S. 2012. Survey on Security Issues in Cloud Computing and Associated Mitigation Techniques. CoRR abs/1204.0764.

Calero, J. A., Edwards, N., Kirschnick, J., Wilcock, L., and Wray, M. 2010. Toward a Multi-Tenancy Authorization System for Cloud Services . *IEEE Security & Privacy 8,* 6 (Nov.-Dec.), 48–55.

Chen, H. and Lee, P. 2014. Enabling Data Integrity Protection in Regenerating-Coding-Based Cloud Storage: Theory and Implementation. *IEEE Transactions on Parallel and Distributed Systems 25,* 2, 407–416.

Chow, R., Golle, P., Jakobsson, M., Shi, E., Staddon, J., Masuoka, R., and Molina, J. 2009. Controlling Data in the Cloud: Outsourcing Computation without Outsourcing Control. In *Proc. of the 2009 ACM Workshop on Cloud Computing Security*. 85–90.

Chu, C., Chow, S., Tzeng, W., Zhou, J., and Deng, R. 2014. Key-Aggregate Cryptosystem for Scalable Data Sharing in Cloud Storage. *IEEE Transactions on Parallel and Distributed Systems 25,* 2, 468–477.

CLOUD-SECURITY-ALLIANCE. 2013. The Notorious Nine: Cloud Computing Top Threats in 2013.

Dabrowski, C. and Mills, K. 2011. VM Leakage and Orphan Control in Open-Source Clouds. In *Proc.* 3^{rd} *IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. 554–559.

Erway, C., Kupcu, A., Papamanthou, C., and Tamassia, R. 2009. Dynamic Provable Data Possession. In *Proceedings of the* 16^{th} *ACM Conference on Computer and Communications Security*.

Factor, M., Hadas, D., Hamama, A., Har'El, N., Kolodner, E., Kurmus, A., and Shulman-Peleg, A. 2013. Secure Logical Isolation for Multi-tenancy in Cloud Storage. In *Proc. MSST*. 1–5.

Feldman, A., Zeller, W., Freedman, M., and Felten, E. 2010. SPORC: Group Collaboration using Untrusted Cloud Resources. In *Proc. OSDI*. 337–350.

Feng, Y., Li, B., and Li, B. 2014. Price Competition in an Oligopoly Market with Multiple IaaS Cloud Providers. *IEEE Transactions on Computers 63,* 1 (January), 59–73.

Godfrey, M. and Zulkernine, M. 2014. Preventing Cache-Based Side-Channel Attacks in a Cloud Environment. *IEEE Transactions on Cloud Computing 2,* 4 (Oct.-Dec.), 395–408.

Gonzalez, N., Miers, C., RedALgolo, F., Carvalho, T., SimplALcio, M., Naslund, M., and Pourzandi, M. 2011. A Quantitative Analysis of Current Security Concerns and Solutions for Cloud Computing. In *Proc. of 3rd IEEE CloudCom*.

Grobauer, B., Walloschek, T., and Stocker, E. 2011. Understanding Cloud Computing Vulnerabilities. *IEEE Security & Privacy 9,* 2 (March-April), 50–57.

Hackett, M. and Hawkey, K. 2012. Security, Privacy and Usability Requirements for Federated Identity . In *In Proceedings of the Workshop on Web 2.0 Security & Privacy*.

He, H., Li, R., Dong, X., and Zhang, Z. 2014. Secure, Efficient and Fine-Grained Data Access Control Mechanism for P2P Storage Cloud. *IEEE Transactions on Cloud Computing 2,* 4 (Oct-Dec), 471–484.

Heilig, L. and Voss, S. 2014. A Scientometric Analysis of Cloud Computing Literature. *IEEE Transactions on Cloud Computing 2,* 3 (July-Sept.), 266–278.

Hogben, G. 2011. ENISA, Cloud Computing: Benefits, Risks and Recommendations for Information Security. https://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment.

Hong, H.-J., Chen, D.-Y., Huang, C.-Y., and Kuan-Ta. 2015. Placing Virtual Machines to Optimize Cloud Gaming Experience. *IEEE Transactions on Cloud Computing 3,* 1 (Jan.- March), 42–53.

Hydara, I., Bakar, A., Sultan, M., Zulzalil, H., and Admodisastro, N. 2015. Current State of Research on Cross-site Scripting (XSS) - A Systematic Literature Review. *Information & Software Technology 58*, 170–186.

Ibrahim, A., Hamlyn-Harris, J., and Grundy, J. 2010. Emerging Security Challenges of Cloud Virtual Infrastructure. In *Proc. of APSEC Cloud Workshop*.

II, J. R. and Al-Hamdani, W. 2011. Who Can You Trust in the Cloud? A Review of Security Issues Within Cloud Computing. In *Proc. of the Information Security Curriculum Development Conference*. 15–19.

Jamshidi, P., Ahmad, A., and Pahl, C. 2013. Cloud Migration Research: A Systematic Review. *IEEE Transactions on Cloud Computing 1,* 2 (July-December), 142–157.

Jansen, W. 2011. Cloud Hooks: Security and Privacy Issues in Cloud Computing. In *Proc.* 44^{th} *Hawaii International Conference on Systems Science*.

Jansen, W. and Grance, T. 2011. Guidelines on Security and Privacy in Public Cloud Computing Special Publication. http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf.

Jayalath, C., Stephen, J., and Eugster, P. 2014. Universal Cross-Cloud Communication. *IEEE Transactions on Cloud Computing 2,* 2 (April-June), 103–116.

Jensen, M., Gruschka, N., and Herkenhner, R. 2009. A Survey of Attacks on Web Services. *Computer Science - R&D 24,* 4 (Nov.), 185–197.

Jensen, M., Schwenk, J., Gruschka, N., and Iacon, L. 2009. On Technical Security Issues in Cloud Computing. In *Proc. of IEEE International Conference on Cloud Computing*. 109–116.

Khalil, I. M., Khreishah, A., and Azeem, M. 2014. Cloud Computing Security: A Survey. *Computers 3,* 1 (Mar.), 1–35.

Kouchaksaraei, H. R. and Chefranov, A. G. 2013. Countering Wrapping Attack on XML Signature in SOAP Message for Cloud Computing . *CoRR abs/1310.0441*.

Le, M. and Tamir, Y. 2011. ReHype: Enabling VM Survival Across Hypervisor Failures. In *Proc. VEE*. 63–74.

Li, J., Li, J., Chen, X., Jia, C., and Lou, W. 2015. Identity-Based Encryption with Outsourced Revocation in Cloud Computing. *IEEE Transactions on Computers 64,* 2 (Feb.), 425–437.

Liu, A., Yuan, Y., and Stavrou, A. 2009. QLProb: A Proxybased Architecture towards Preventing SQL Injection Attacks. In *Proc. SAC*.

Liu, X., Zhang, Y., Wang, B., and Yan, J. 2013. Mona: Secure Multi-Owner Data Sharing for Dynamic Groups in the Cloud. *IEEE Transactions on Parallel and Distributed Systems 24,* 6, 1182–1191.

Lombardi, F. and Pietro, R. D. 2011. Secure Virtualization for Cloud Computing. *Journal of Network and Computer Applications 34,* 4 (July), 1113–1122.

Luo, S., Lin, Z., Chen, X., Yang, Z., and Chen, J. 2011. Virtualization Security for Cloud Computing Services. In *Proc. Int. Conf on Cloud and Service Computing*. 174–179.

Maler, E. and Reed, D. 2008. The Venn of Identity: Options and Issues in Federated Identity Management . *IEEE Security & Privacy 6,* 2 (Apr.), 1623.

Mastroianni, C., Meo, M., and Papuzzo, G. 2013. Probabilistic Consolidation of Virtual Machines in Self-Organizing Cloud Data Centers. *IEEE Transactions on Cloud Computing 1,* 2 (July-December), 215–228.

Mather, T., Kumaraswamy, S., and latif, S. 2009. *Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance - 1^{st} edition.* OReilly Media.

Mell, P. and Grance, T. 2011. The NIST Definition of Cloud Computing - Special Publication 800-145. National Institute of Standards and Technology.

Moreno, I. S., Garraghan, P., Townend, P., and Xu, J. 2014. Analysis, Modeling and Simulation of Workload Patterns in a Large-Scale Utility Cloud. *IEEE Transactions on Cloud Computing 2,* 2 (April-June), 208–221.

Morshedlou, H. and Meybodi, M. R. 2014. Decreasing Impact of SLA Violations:A Proactive Resource Allocation Approachfor Cloud Computing Environments. *IEEE Transactions on Cloud Computing 2,* 2 (April-June), 156–167.

Morsy, M., Grundy, J., and Mller, I. 2010. An Analysis of the Cloud Computing Security Problem. In *Proc. of APSEC Cloud Workshop*.

Mundada, Y., Ramachandran, A., and Feamster, N. 2011. SilverLine: Data and Network Isolation for Cloud Services. In *Proc. HotCloud*.

Papagianni, C., Leivadeas, A., Papavassiliou, S., Maglaris, V., and Monje, C. C.-P. A. 2013. On the Optimal Allocation of Virtual Resources in Cloud Computing Networks. *IEEE Transactions on Computers 62,* 6 (June), 1060–1071.

Pearce, M., Zeadally, S., and Hunt, R. 2013. Virtualization: Issues, security threats, and solutions . *ACM Comput. Surv. 45,* 2 (Apr.), 17.

Pearson, S. and Benameur, A. 2010. Privacy, Security and Trust Issues Arising from Cloud Computing . In *Proc. of IEEE 2^{nd} International Conference on Cloud Computing Technology and Science (CloudCom)*.

Prasad, A. S. and Rao, S. 2014. A Mechanism Design Approach to Resource Procurement in Cloud Computing. *IEEE Transactions on Computers 63,* 1 (January), 17–30.

Rahumed, A., Chen, H., Tang, Y., Lee, P., and Lui, J. 2011. A Secure Cloud Backup System with Assured Deletion and Version Control. In *Proc. 3^{rd} Intl Workshop Security in Cloud Computing*.

Reimer, D., Thomas, A., Ammons, G., Mummert, T. W., Alpern, B., and Bala, V. 2008. Opening Black Boxes: Using Semantic Information to Combat Virtual Machine Image Sprawl. In *Proc. VEE*. 111–120.

Ristenpart, T., Tromer, E., Shacham, H., and Savage, S. 2009. Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds. In *Proc. ACM Conference on Computer and Communications Security*.

Ruj, S., Stojmenovic, M., and Nayak, A. 2014. Decentralized Access Control with Anonymous Authentication of Data Stored in Clouds. *IEEE Transactions on Parallel and Distributed Systems 25,* 2 (Feb.).

Sabahi, F. 2012. Secure Virtualization for Cloud Environment Using Hypervisor-based Technology. In *Proc. Int. Journal of Machine Learning and Computing*. Vol.2, No.2, 39–45.

Schwarzkopf, R., Schmidt, M., Fallenbeck, N., and Freisleben, B. 2009. Multi-layered Virtual Machines for Security Updates in Grid Environments. In *Proc. EUROMICRO-SEAA*. 563–570.

Schwarzkopf, R., Schmidt, M., Fallenbeck, N., and Freisleben, B. 2011. Checking Running and Dormant Virtual Machines for the Necessity of Security Updates in Cloud Environments. In *Proc. CloudCom*. 239–246.

Schwarzkopf, R., Schmidt, M., Strack, C., Martin, S., and Freisleben, B. 2012. Increasing Virtual Machine Security in Cloud Environments. *Journal of Cloud Computing*.

Sosinsky, B. 2011. *Cloud Computing Bible*. Wiley Publications.

Spring, J. 2011. Monitoring Cloud Computing by Layer, Part 1. *IEEE Security & Privacy 9,* 2 (March-April), 66–68.

Subashini, S. and V.Kavitha. 2011. A Survey on Security Issues in Service Delivery Models of Cloud Computing. *Journal of Network and Computer Applications 34,* 1 (January), 1–11.

Takabi, H., Joshi, J., and Ahn, G. 2010. Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security & Privacy 8,* 6, 24–31.

Tang, Y., Lee, P., Lui, J., and Perlman, R. 2010. FADE: Secure Overlay Cloud Storage with File Assured Deletion. In *Proc.* 6^{th} *Intl ICST Conf. Security and Privacy in Comm. Networks (SecureComm)*.

Tysowski, P. K. and Hasan, M. A. 2013. Hybrid Attribute- and Re-Encryption-Based Key Management for Secure and Scalable Mobile Applications in Clouds. *IEEE Transactions on Cloud Computing 1,* 2 (July-December), 172–186.

Wang, B., Li, B., and Li, H. 2014. Oruta: Privacy-preserving Public Auditing for Shared Data in the Cloud. *IEEE Transactions on Cloud Computing 2,* 1 (Jan. - March), 43–56.

Wang, C., Chow, S., Wang, Q., Ren, K., and Lou, W. 2011. Privacy-Preserving Public Auditing for Secure Cloud Storage. *IEEE Transactions on Computers 62,* 2 (Feb.), 362–375.

Wang, C., Wang, Q., Ren, K., and Lou, W. 2009. Ensuring Data Storage Security in Cloud Computing. In *Proceedings of the* 17^{th} *International Workshop on Quality of Service*. 1–9.

Wang, Q., Wang, C., Li, J., Ren, K., and Lou, W. 2009. Enabling Public Verifiability and Data Dynamics for Storage Security. In *Proceedings of the* 14^{th} *European Conference on Research in Computer Security*.

Wang, Q., Wang, C., Ren, K., Lou, W., and Li, J. 2011. Enabling Public Auditability and Data Dynamics for Storage Security in Cloud Computing. *IEEE Transactions on Parallel and Distributed Systems 22,* 5 (May).

Wang, W., Li, Z., Owens, R., and Bhargava, B. 2009. Secure and Efficient Access to Outsourced Data. In *Proc. ACM Workshop Cloud Computing Security (CCSW)*.

Wang, Y. and Shi, W. 2014. Budget-Driven Scheduling Algorithms for Batches of MapReduce Jobs in Heterogeneous Clouds. *IEEE Transactions on Cloud Computing 2,* 3 (July-Sept.), 306–319.

Wei, J., Zhang, X., Ammons, G., Bala, V., and Ning, P. 2009. Managing Security of Virtual Machine Images in a Cloud Environment. In *Proc. CCSW*. 91–96.

Xue, K. and Hong, P. 2014. A Dynamic Secure Group Sharing Framework in Public Cloud Computing. *IEEE Transactions on Cloud Computing 2,* 4 (Oct.-Dec), 459–470.

Yan, L., Rong, C., and Zhao, G. 2009. Strengthen Cloud Computing Security with Federal Identity Management Using Hierarchical Identity-Based Cryptography. In *Proceedings of the* 1^{st} *International Conference on Cloud Computing*. 167–177.

Yang, K. and Jia, X. 2013. An Efficient and Secure Dynamic Auditing Protocol for Data Storage in Cloud Computing. *IEEE Transactions on Parallel and Distributed Systems 24,* 9, 1717–1726.

Yu, S., Wang, C., Ren, K., and Lou, W. 2010. Achieving Secure, Scalable, and Fine-grained Data Access Control in Cloud Computing. In *Proceedings of IEEE INFOCOM*.

Zaman, S. and Grosu, D. 2013. A Combinatorial Auction-Based Mechanism for Dynamic VM Provisioning and Allocation in Clouds. *IEEE Transactions on Cloud Computing 1,* 2 (July-December), 129–141.

Zeng, K. 2008. Publicly Verifiable Remote Data Integrity. In *Proceedings of the* 10^{th} *International Conference on Information and Communications Security*. 419–434.

Zhang, Y., Ion, M., Russello, G., and Crispo, B. 2012. Cross-VM Side Channels and their Use to Extract Private Keys. In *Proc. ACM CCCS*. 305–316.

Zissis, D. and Lekkas, D. 2012. Addressing Cloud Computing Security Issues. *Future Generation Comp. Syst. 28,* 3, 583–592.

Zunnurhain, K. and Vrbsky, S. 2010. Security Attacks and Solutions in Clouds. In *Proc.* 1^{st} *International Conference on Cloud Computing*. 145156.

Zunnurhain, K., Vrbsky, S. V., and Hasan, R. 2014. FAPA: Flooding Attack Protection Architecture in a Cloud System. *International Journal of Cloud Computing 3,* 4 (Nov.), 379–401.

Mr. Tariqul Islam received his B.S degree in Computer Science and Engineering from University of Dhaka, Bangladesh, in 2008. He is currently a Ph.D student in the Department of Computer Science at University of Kentucky. His research interest lies in issues related to security, privacy and trust in cloud computing.

Dr. D. Manivannan is currently an associate professor of Computer Science at University of Kentucky, Lexington, Kentucky, USA. He received an M.Sc degree in mathematics from University of Madras, Madras, India. He received M.S and PhD degrees in computer and information science from The Ohio State University, Ohio, in 1993 and 1997 respectively. He published his research work in the following areas: fault-tolerance and synchronization in distributed systems, routing in wormhole networks, routing in ad hoc networks and vehicular ad hoc networks, channel allocation in cellular networks, wireless personal area networks and sensor networks. Dr. Manivannan has published more than 60 articles in refereed International Journals (a vast majority of which were published by IEEE, ACM, Elsevier, and Springer) and Proceedings of International Conferences.
He served as an Associate Editor of IEEE Transactions on Parallel and Distributed Systems and IEEE Communications Magazine. He served as Program co-chair of three International Conferences in the areas of reliable distributed systems and wireless networks and served as program committee member for over 40 International Conferences. He is on the Editorial Board of Information Sciences journal and Wireless Personal Communications journal. He served as reviewer for more than 30 International Journals published by ACM, IEEE, Elsevier, Springer, Oxford University Press, Taylor and Francis and others. He also served on several proposal review panels of US National Science Foundation and as external tenure reviewer for other Universities. Dr. Manivannan's research has been funded by grants from the US National Science Foundation and the US Department of Treasury.
Dr. Manivannan is a recipient of the Faculty Early Career Development award (CAREER award) from US National Science Foundation. He is a senior member of the IEEE and ACM.

Dr. Sherali Zeadally received his Bachelor's degree in computer science from the University of Cambridge, United Kingdom, and his doctoral degree in computer science from the University of Buckingham, United Kingdom. He is an associate professor in the College of Communication and Information at the University of Kentucky. He is a Fellow of the British Computer Society and a Fellow of the Institution of Engineering Technology, United Kingdom.