Requirements and Challenges for Securing Cloud Applications and Services

Mrs. Y. Lakshmi Prasanna[1], Mrs. S. Neelima[2], Mrs. M. Padmavathi[3]
Department of Information Technology[1], Department of Master of Computer Applications[2], Department of Master of Computer Applications[3]
Swarna Bharathi Institute of Science and Technology, Pakabanda Bazaar, Khammam, Andhra Pradesh.

ABSTRACT: *Adopting cloud computing is a complex decision involving many factors. This paper focuses on the central issues of cloud computing security. The Cloud Computing Architectural Framework, provides a conceptual framework to help evaluate initial cloud risks and inform security decisions. It's a quick method for evaluating the tolerance for moving an asset to various cloud computing models. The keys to understanding how cloud architecture impacts security architecture are a common and concise lexicon, coupled with a consistent taxonomy of offerings by which cloud services and architecture can be deconstructed, mapped to a model of compensating security and operational controls, risk assessment and management frameworks, and in turn to compliance standards. Understanding the relationships and dependencies between Cloud computing models is critical to understanding Cloud computing security risks. The purpose of this paper is to provide needed context for assistance in making educated risk management decisions regarding their cloud adoption strategies.*

Keywords – *Cloud Architecture, Cloud Operations, Cloud Security, Reference Models, Risk Management.*

I. INTRODUCTION

Cloud Computing is still a rapidly evolving landscape; and one that requires us to stay current or fall behind. As the march of Cloud Computing continues, it brings both new opportunities and new security challenges. Cloud Computing isn't necessarily more or less secure; as with any new technology, it creates new risks and new opportunities. In some cases moving to the cloud provides an opportunity to re-architect older applications and infrastructure to meet or exceed modern security requirements. At other times the risk of moving sensitive data and applications to an emerging infrastructure might exceed the tolerance. Main goal is to provide with practical recommendations and key points to make that transition as securely as possible.

Cloud computing ('cloud') is an evolving term that describes the development of many existing technologies and approaches to computing into something different. Cloud separates application and information resources from the underlying infrastructure, and the mechanisms used to deliver them.

Cloud enhances collaboration, agility, scaling, and availability, and provides the potential for cost reduction through optimized and efficient computing. More specifically, cloud describes the use of a collection of services, applications, information, and infrastructure comprised of pools of compute, network, information, and storage resources. These components can be rapidly orchestrated, provisioned, implemented and decommissioned, and scaled up or down; providing for an on-demand utility-like model of allocation and consumption.

With so many different cloud deployment options — including the SPI service models; public vs. private deployments, internal vs. external hosting, and various hybrid permutations — no list of security controls can cover all circumstances. As with any security area, organizations should adopt a risk-based approach to moving to the cloud and selecting security options.

From an architectural perspective; there is much confusion surrounding how cloud is both similar to and different from existing models of computing; and how these similarities and differences impact the organizational, operational, and technological approaches to network and information security practices.

The keys to understanding how cloud architecture impacts security architecture are a common and concise lexicon, coupled with a consistent taxonomy of offerings by which cloud services and architecture can be deconstructed, mapped to a model of compensating security and operational controls, risk assessment and management frameworks, and in turn to compliance standards.

II. CLOUD ARCHITECTURE

The Cloud Computing Architectural Framework helps in evaluating the initial cloud risks and to inform security decisions. This is a simple framework defining many of the concepts and domains as follows:

2.1 Identify the asset for the cloud deployment

At the simplest, assets supported by the cloud fall into two general buckets:

2.1.1. Data

2.1.2. Applications/Functions/Processes

With cloud computing data and applications don't need to reside in the same location, and can even shift only parts of functions to the cloud. For example, we can host our application and data in our own data center, while still outsourcing a portion of its functionality to the cloud through a Platform as a Service.

The first step in evaluating risk for the cloud is to determine exactly what data or function is being considered for the cloud. This should include potential uses of the asset once it moves to the cloud to account for scope creep. Data and transaction volumes are often higher than expected.

2.2 Evaluate the asset

The next step is to determine how important the data or function is to the organization. This can be assessed by considering the factors confidentiality, integrity, and availability requirements for the asset; and how those are affected if all or part of the asset is handled in the cloud.

2.3 Map the asset to potential cloud deployment models

The next step is to determine which deployment models are comfortable with and also looking at potential providers to that can accept the risks implicit to the various deployment models: private, public, community, or hybrid; and hosting scenarios: internal, external, or combined.

2.4 Evaluate potential cloud service models and providers

In this step focus is on the degree of control having at each SPI tier to implement any required risk management and on the degree of control to implement risk mitigations in the different SPI tiers.

III. VISUAL MODEL OF CLOUD COMPUTING

NIST defines cloud computing by describing five essential characteristics, three cloud service models, and four cloud deployment models. They are summarized in visual form in figure 1 and explained in detail below.

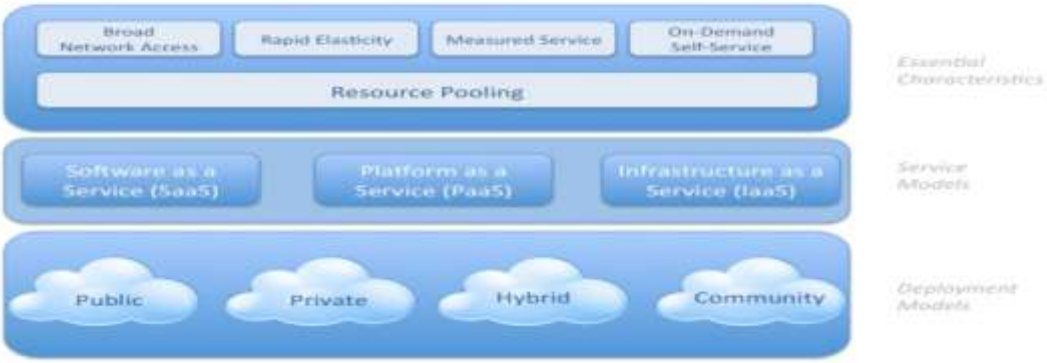

Fig 1. NIST Visual Model of Cloud Computing Definition

In the short term, as innovation drives rapid solution development, consumers and providers of cloud services will enjoy varied methods of interacting with cloud services in the form of developing APIs and interfaces and so cloud service brokers will emerge as an important component in the overall cloud ecosystem.

Cloud service brokers will abstract these possibly incompatible capabilities and interfaces on behalf of consumers to provide proxy in advance of the arrival of common, open and standardized ways of solving the problem longer term with a semantic capability that allows fluidity and agility in a consumer being able to take advantage of the model that works best for their particular needs.

It is also important to note that there are derivative cloud deployment models emerging due to the maturation of market offerings and customer demand. An example of such is virtual private clouds — a way of utilizing public cloud infrastructure in a private or semi-private manner and interconnecting these resources to the internal resources of a consumers' datacenter, usually via virtual private network (VPN) connectivity.

3.1 Multi-tenancy

Although multi-tenancy is not considered as an essential characteristic of Cloud Computing in NIST's model, it is an important element of cloud. Multi-tenancy in cloud service models implies a need for policy-driven enforcement, segmentation, isolation, governance, service levels, and chargeback/billing models for different consumer constituencies. Consumers might utilize a public cloud provider's service offerings or actually be from the same organization, such as different business units rather than distinct organizational entities, but would still share infrastructure.

From a provider's perspective, multi-tenancy suggests an architectural and design approach to enable economies of scale, availability, management, segmentation, isolation, and operational efficiency; leveraging shared infrastructure, data, metadata, services, and applications across many different consumers.

Multi-tenancy can also take on different definitions depending upon the cloud service model of the provider; in as much as it may entail enabling the capabilities described above at the infrastructure, database, or application levels. An example would be the difference between an IaaS and SaaS multi-tenant implementation.

Cloud deployment models place different importance on multi-tenancy. However, even in the case of a private cloud, a single organization may have a multitude of third party consultants and contractors, as well as a desire for a high degree of logical separation between business units. Thus multi-tenancy concerns should always be considered.

3.2 Cloud Reference Model

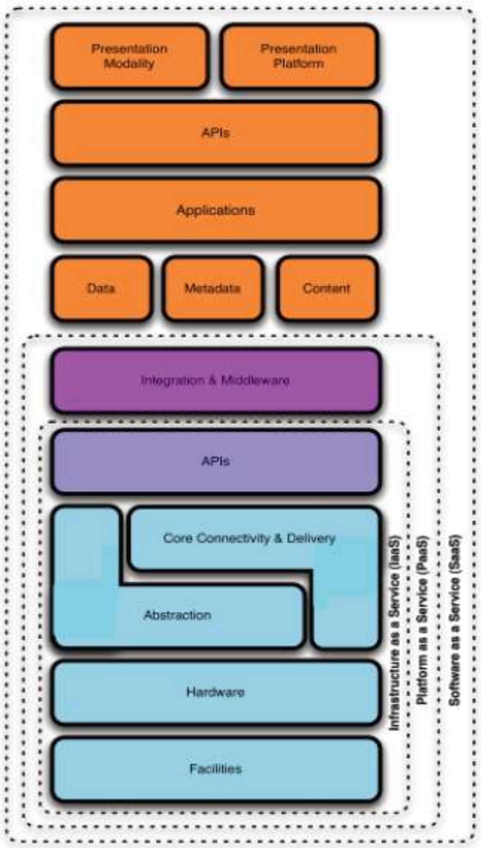

Fig 2. Cloud Reference Model

Understanding the relationships and dependencies between Cloud Computing models is critical to understanding Cloud Computing security risks. IaaS is the foundation of all cloud services, with PaaS building upon IaaS, and SaaS in turn building upon PaaS as described in the Cloud Reference Model diagram. In this way, just as capabilities are inherited, so are information security issues and risk.

IaaS includes the entire infrastructure resource stack from the facilities to the hardware platforms that reside in them. It incorporates the capability to abstract resources (or not), as well as deliver physical and logical connectivity to those resources. Ultimately, IaaS provides a set of APIs which allow management and other forms of interaction with the infrastructure by consumers.

PaaS sits atop IaaS and adds an additional layer of integration with application development frameworks; middleware capabilities; and functions such as database, messaging, and queuing; which allow developers to build applications upon to the platform; and whose programming languages and tools are supported by the stack.

SaaS in turn is built upon the underlying IaaS and PaaS stacks; and provides a self-contained operating environment used to deliver the entire user experience including the content, its presentation, the application(s), and management capabilities.

It should therefore be clear that there are significant trade-offs to each model in terms of integrated features, complexity vs. openness (extensibility), and security. Trade-offs between the three cloud deployment models include:

- Generally, SaaS provides the most integrated functionality built directly into the offering, with the least consumer extensibility, and a relatively high level of integrated security (at least the provider bears a responsibility for security).
- PaaS is intended to enable developers to build their own applications on top of the platform. As a result it tends to be more extensible than SaaS, at the expense of customer ready features. This tradeoff extends to security features and capabilities, where the built-in capabilities are less complete, but there is more flexibility to layer on additional security.

- IaaS provides few if any application-like features, but enormous extensibility. This generally means less integrated security capabilities and functionality beyond protecting the infrastructure itself. This model requires that operating systems, applications, and content be managed and secured by the cloud consumer.

The key take-away for security architecture is that the lower down the stack the cloud service provider stops, the more security capabilities and management consumers are responsible for implementing and managing themselves.

In the case of SaaS, this means that service levels, security, governance, compliance, and liability expectations of the service and provider are contractually stipulated; managed to; and enforced. In the case of PaaS or IaaS it is the responsibility of the consumer's system administrators to effectively manage the same, with some offset expected by the provider for securing the underlying platform and infrastructure components to ensure basic service availability and security. It should be clear in either case that one can assign/transfer responsibility but not necessarily accountability.

Narrowing the scope or specific capabilities and functionality within each of the cloud delivery models, or employing the functional coupling of services and capabilities across them, may yield derivative classifications. For example "Storage as a Service" is a specific sub-offering within the IaaS 'family'.

While a broader review of the growing set of cloud computing solutions is outside the scope of this document, the OpenCrowd Cloud Solutions taxonomy in the figure below provides an excellent starting point. The OpenCrowd taxonomy demonstrates the swelling ranks of solutions available today across each of the previously defined models.

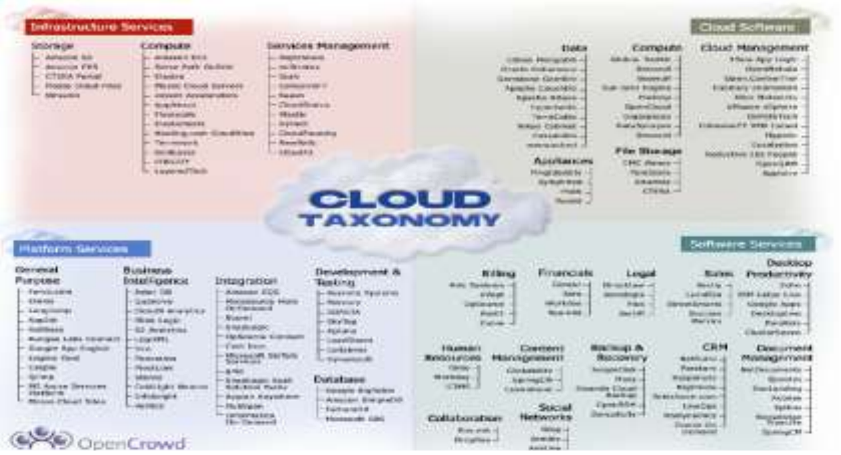

Fig 3. OpenCrowd Taxonomy

To describe and define common cases and demonstrate the benefits of cloud, with the goal being to "...bring together cloud consumers and cloud vendors to define common use cases for cloud computing...and highlight the capabilities and requirements that need to be standardized in a cloud environment to ensure interoperability, ease of integration, and portability."

IV. CLOUD SECURITY REFERENCE MODEL

The cloud security reference model addresses the relationships of these classes and places them in context with their relevant security controls and concerns. The deployment and consumption modalities of cloud should be thought of not only within the context of 'internal' vs. 'external' as they relate to the physical location of assets, resources, and information; but also by whom they are being consumed by; and who is responsible for their governance, security, and compliance with policies and standards.

This is not to suggest that the on- or off-premise location of an asset, a resource, or information does not affect the security and risk posture of an organization because they do — but to underscore that risk also depends upon:

- The types of assets, resources, and information being managed
- Who manages them and how
- Which controls are selected and how they are integrated
- Compliance issues

For example

- ✓ a LAMP stack deployed on Amazon's AWS EC2 would be classified as a public, off-premise, third-party managed-IaaS solution; even if the instances and applications/data contained within them were managed by the consumer or a third party.
- ✓ A custom application stack serving multiple business units; deployed on Eucalyptus under a corporation's control, management, and ownership; could be described as a private, on-premise, self-managed SaaS solution. Both examples utilize the elastic scaling and self-service capabilities of cloud.

Another way of visualizing how combinations of cloud service models, deployment models, physical locations of resources, and attribution of management and ownership, is the Jericho Forum's Cloud Cube Model, shown in the figure below:

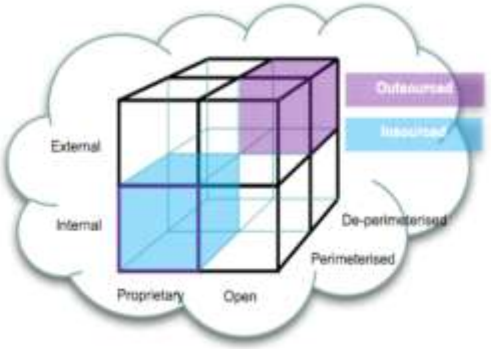

Fig 4. Jericho Forum's Cloud Cube Model,

The Cloud Cube Model illustrates the many permutations available in cloud offerings today and presents four criteria/dimensions in order to differentiate cloud 'formations' from one another and the manner of their provision, in order to understand how cloud computing affects the way in which security might be approached.

The Cloud Cube Model also highlights the challenges of understanding and mapping cloud models to control frameworks and standards such as ISO/IEC 27002, which provides "...a series of guidelines and general principles for initiating, implementing, maintaining, and improving information security management within an organization."

V. MAPPING THE CLOUD MODEL TO SECURITY CONTROLS & COMPLIANCE MODEL

The figure below shows an example of how a cloud service mapping can be compared against a catalogue of compensating controls to determine which controls exist and which do not — as provided by the consumer, the cloud service provider, or a third party. This can in turn be compared to a compliance framework or set of requirements such as PCI DSS, as shown.

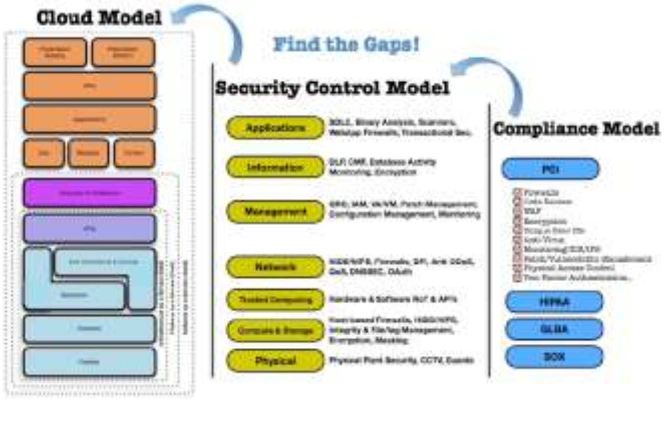
Cloud Model
Find the Gaps!
Security Control Model
Applications
SDLC, Binary Analysis, Scanners, WebApp Firewalls, Transactional Sec.
Information
DLP, CMF, Database Activity Monitoring, Encryption
Management
GRC, IAM, VA/VM, Patch Management, Configuration Management, Monitoring
Network
NIDS/NIPS, Firewalls, DPI, Anti-DDoS, QoS, DNSSEC, OAuth
Trusted Computing
Hardware & Software RoT & API's
Compute & Storage
Host-based Firewalls, HIDS/HIPS, Integrity & File/log Management, Encryption, Masking
Physical
Physical Plant Security, CCTV, Guards
Compliance Model
PCI
HIPAA
GLBA
SOX

Fig 5.Mapping the Cloud Model to the Security Control & Compliance Model

Once this gap analysis is complete, per the requirements of any regulatory or other compliance mandates, it becomes much easier to determine what needs to be done in order to feed back into a risk assessment framework; this, in turn, helps to determine how the gaps and ultimately risk should be addressed: accepted, transferred, or mitigated.

It is important to note that the use of cloud computing as an operational model does not inherently provide for or prevent achieving compliance. The ability to comply with any requirement is a direct result of the service and deployment model utilized and the design, deployment, and management of the resources in scope.

VI. SECURITY FOR CLOUD COMPUTING

Security controls in cloud computing are, for the most part, no different than security controls in any IT environment. However, because of the cloud service models employed, the operational models, and the technologies used to enable cloud services, cloud computing may present different risks to an organization than traditional IT solutions.

Cloud computing is about gracefully losing control while maintaining accountability even if the operational responsibility falls upon one or more third parties.

An organization's security posture is characterized by the maturity, effectiveness, and completeness of the risk-adjusted security controls implemented. These controls are implemented in one or more layers ranging from the facilities (physical security), to the network infrastructure (network security), to the IT systems (system security), and all the way to the information and applications (application security). Additionally controls are implemented at the people and process levels, such as separation of duties and change management, respectively.

As described earlier in this document, the security responsibilities of both the provider and the consumer greatly differ between cloud service models. Amazon's AWS EC2 infrastructure as a service offering, as an example, includes vendor responsibility for security up to the hypervisor, meaning they can only address security controls such as physical security, environmental security, and virtualization security. The consumer, in turn, is responsible for security controls that relate to the IT system (instance) including the operating system, applications, and data.

The inverse is true for Salesforce.com's customer resource management (CRM) SaaS offering. Because the entire 'stack' is provided by Salesforce.com, the provider is not only responsible for the physical and environmental security controls, but it must also address the security controls on the infrastructure, the applications, and the data. This alleviates much of the consumer's direct operational responsibility.

One of the attractions of cloud computing is the cost efficiencies afforded by economies of scale, reuse, and standardization. To bring these efficiencies to bear, cloud providers have to provide services that are flexible enough to serve the largest customer base possible, maximizing their addressable market. Unfortunately, integrating security into these solutions is often perceived as making them more rigid.

The figure below illustrates these issues: in SaaS environments the security controls and their scope are negotiated into the contracts for service; service levels, privacy, and compliance are all issues to be dealt with legally in contracts. In an IaaS offering, while the responsibility for securing the underlying infrastructure and abstraction layers belongs to the provider, the remainder of the stack is the consumer's responsibility. PaaS offers a balance somewhere in between, where securing the platform itself falls onto the provider, but securing the applications developed against the platform and developing them securely, both belong to the consumer.

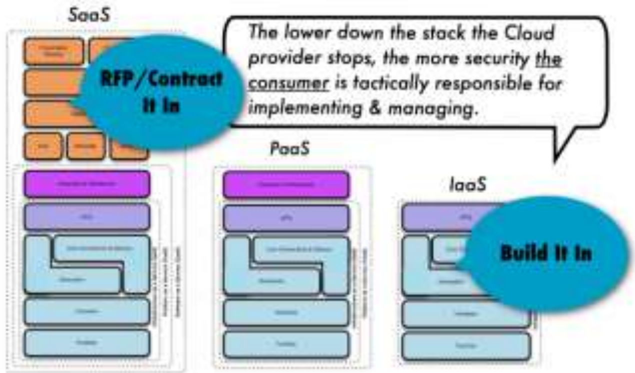
SaaS
RFP/Contract It In
The lower down the stack the Cloud provider stops, the more security the consumer is tactically responsible for implementing & managing.
PaaS
IaaS
Build It In

Fig 6. Security Integration into SPI service models

VII. CONCLUSION

The importance of assets moving to the cloud is considered and combinations of deployment and service models acceptable are presented. A rough idea of potential exposure points for sensitive information and operations is also made available. These together gives sufficient context needed to evaluate security controls.

For low-value assets the same level of security controls can be applied and can skip many of the recommendations — such as on-site inspections, discoverability, and complex encryption schemes. A high-value regulated asset might entail audit and data retention requirements.

Due to our limited space, as well as the depth and breadth of material to cover, this document contains extensive lists of security recommendations. Not all cloud deployments need every possible security and risk control. Spending a little time up front evaluating risk tolerance and potential exposures will provide the context needed to pick and choose the best options for an organization and deployment.

The keys to understanding how cloud architecture impacts security architecture are coupled with a consistent taxonomy of offerings by which cloud services and architecture can be deconstructed, mapped to a model of compensating security and operational controls, risk assessment frameworks, and management frameworks and in turn to compliance standards.

Understanding how architecture, technology and process change or remain the same when deploying Cloud Computing services is critical. Without a clear understanding of the higher-level architectural implications, it is impossible to address more detailed issues rationally.

This architectural overview will provide a solid foundation for assessing, operationalizing, managing, and governing security in Cloud Computing environments.

ACKNOWLEDGEMENT

The authors would like to acknowledge and express their heartfelt gratitude to the Swarna Bharathi Institute of Science and Technology [SBIT] Management, Chairman, Director, Principal and Dean R&D. for their vital encouragement, support and facilities providing by them to achieve this task and at they would like to thank Anonymous Reviewers for their valuable suggestions and comments. This paper has greatly benefited from their Efforts.

REFERENCES

[1] Security Guidance for Critical Areas of Focus in Cloud Computing, Version 1, by Cloud Security Alliance, April 2009

[2] Amazon web services blog: Introducing amazon virtual private cloud (vpc), Amazon, August 2009. Amazon Web Services: Overview of Security Processes, September 2008

[3] A guide to security metrics. SANS Institute, June 2006. http://www.sans.org

[4] Amazon EC2 API - http://docs.amazonwebservices.com/AWSEC2/2006-10-01/DeveloperGuide/

[5] Amazon Elastic Compute Cloud Developer Guide, http://docs.amazonwebservices.com/AWSEC2/2009-03-01/DeveloperGuide/

[6] Business Software Alliance, Information Security Governance: Towards a Framework for Action Centers for Medicare and Medicaid Services Information Security Risk Assessment Methodology Cloud Computing and Compliance: Be Careful Up There (Wood, Lamont, ITWorld, January 30, 2009)

[7] Cloud computing definition, by P. Mell and T. Grance, NIST June 2009. http://csrc.nist.gov/groups/SNS/cloud-computing/index.html

[8] Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration, Jericho Forum, V 1.0, April 2009

[9] Cloud Computing and Compliance: Be Careful Up There, Wood, Lamont, ITWorld, January 30, 2009

[10] Jericho Forum - http://www.opengroup.org/jericho/ and the Jericho Cloud Cube model - http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

[11] NIST Recommended Security Controls for Federal Information Systems (SP800-53)

[12] NIST SP 800-30 Risk Management Guide for Information Technology Systems

[13] Open Cloud Computing Interface Working Group - http://www.occi-wg.org/doku.php

[14] Open Security Architecture Group - http://www.opensecurityarchitecture.org

[15] OpenCrowd - http://www.opencrowd.com/views/cloud.php Security Guidance for Critical Areas of Focus in Cloud Computing, Version 1, by Cloud Security Alliance, April 2009

Mrs. Y. Lakshmi Prasanna, working as an Associate professor in the department of Information Technology. Her research areas include Cloud Computing, Network Security, Computer Networks and Mobile Computing.

Mrs. S. Neelima, working as an Associate professor in the department of Master of Computer Applications. Her research areas include Cloud Computing, Data Ware housing and data mining, Computer Networks and Network Security.

Mrs. M. Padmavathi, working as an Associate professor in the Department of Master of Computer Applications. Her research areas include Cloud Computing, Mobile Computing, Data Ware housing and data mining and Image Processing.