An Analysis of Security Issues and Protections Against vulnerabilities in Cloud Computing – A Review

[1]Dr. M E Purushoththaman
Professor & Principal
Swami Vivekanada Institute of Technology
Department of Computer Science and Engineering
Hyderabad, India

[2]B Bhavani
Research Scholar,
Shri Jagadishprasad Jhabarmal Tibrewala University.
Department of Computer Science and Engineering
JhunJhunu, India

Abstract: *Cloud computing is a model to enable existing anywhere and everywhere, easy to use, on-demand access to a shared pool of configurable computing resources like., networks, servers, storage, applications, and other services. It is a disruptive technology with potential to enhance collaboration, agility, scaling, availability and provides the cost reduction through optimized and efficient computing. The cloud model brings insight, a world where components can be rapidly orchestrated, provisioned, implemented and decommissioned, and scaled up or down to provide an on-demand utility-like model of allocation and consumption. There is a thin line between conventional computing and cloud computing in terms of the organizational, operational, and technological approaches to network and information security practices. There are many definitions today that attempt to address cloud from the perspective of academicians, architects, engineers, developers, managers, and consumers. This paper focuses on cloud computing architecture models and their features, advantages and security related issues on every concerned that is specifically tailored to the unique perspectives of IT network and security professionals.*

Keywords: *Cloud Computing, Networks, Servers, Agility, Scaling, Provisioning.*

1. INTRODUCTION

Cloud computing is one of the most significant technological development of our time. It gives the pleasure of utilizing and paying for this utilization in terms of resources and time. Customers, of all ranges, whether or small businesses or large enterprises, are attracted toward the cloud's capability of agility, reduced capital costs, and availability of enhanced IT resources instantly. IT companies are shifting from providing their own IT infrastructure to utilizing the computation services provided by the cloud for their information technology needs (Carr, 2008). Cloud computing ensures a level of abstraction between the physical infrastructure and the owner of the information being stored and processed. Such indirect control of the physical environment introduces vulnerabilities unknown in previous settings. Such a radical change is of course not risk free. As IT services are contracted outside of the enterprise, the dependency on third party providers compels companies to rethink their risk management techniques and adapt accordingly. WE organized this paper is organized as . CLOUD REFERENCE MODEL Security for Cloud Computing, Cloud Deployment Models, Information management and data security, Information Governance, Data Security, Data Loss Prevention, Database and File Activity Monitoring and we are forced to suppress many related topics as the content exceeds more and more.

II. 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 capabilities are inherited, so are information security issues and risks. 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 API's, which allows management and other forms of interaction with the infrastructure by consumers.

- ***Infrastructure as a Service (IaaS),*** delivers computer infrastructure (typically a platform virtualization environment) as a service, along with raw storage and networking. Rather than purchasing servers, software, data-center space, or network equipment, clients instead buy those resources as a fully outsourced service.
- *Software as a service (SaaS),* sometimes referred to as "on-demand software," is a software delivery model in which software and its associated data are hosted centrally (typically in the (Internet) cloud) and are typically accessed by users using a thin client, normally using a web browser over the Internet.
- ***Platform as a service (PaaS),*** is the delivery of a computing platform and solution stack as a service. PaaS offerings facilitate deployment of applications without the cost and complexity of buying and managing the underlying hardware and software and provisioning hosting capabilities. This provides all

of the facilities required to support the complete life cycle of building and delivering web applications and services entirely available from the Internet.
The deployment and consumption modalities of cloud should be thought of not only within the context of 'internal' versus 'external' as they relate to the physical location of assets, resources, and information; but also by whom they are being consumed; 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

The following table summarizes these points:
Table 1—Cloud Computing Deployment Models

	Infrastructure Managed By[1]	Infrastructure Owned By[2]	Infrastructure Located[3]	Accessible and Consumed By[4]
Public	Third Party Provider	Third Party Provider	Off-Premise	Untrusted
Private/ Community	Organization Or Third Party Provider	Organization Third Party Provider	On-Premise Off-Premise	Trusted
Hybrid	Both Organization & Third Party Provider	Both Organization & Third Party Provider	Both On-Premise & Off-Premise	Trusted & Untrusted

[1] Management includes: governance, operations, security, compliance, etc...
[2] Infrastructure implies physical infrastructure such as facilities, compute, network & storage equipment
[3] Infrastructure Location is both physical and relative to an Organization's management umbrella and speaks to ownership versus control
[4] Trusted consumers of service are those who are considered part of an organization's legal/contractual/ policy umbrella including employees, contractors, & business partners. Untrusted consumers are those that may be authorized to consume some/all services but are not logical extensions of the organization.

III. 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.

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), 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 Salesforce.com provides the entire "stack," 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.

There is currently no way for a naive consumer of cloud services to simply understand what exactly he/she is responsible for [though reading this guidance document should help], but there are efforts underway by the CSA and other bodies to define standards around cloud audit.

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.

This rigidity often manifests in the inability to gain parity in security control deployment in cloud environments compared to traditional IT. This stems mostly from the

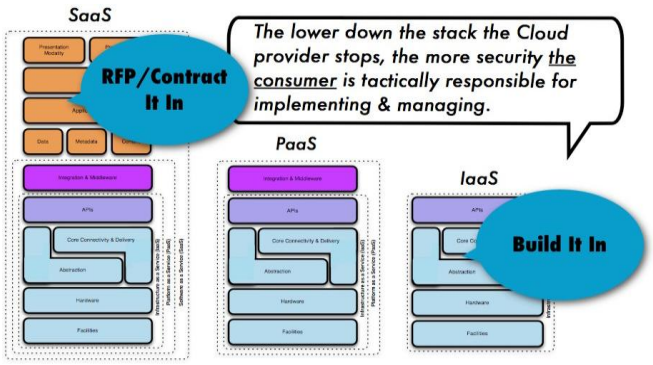

abstraction of infrastructure, and the lack of visibility and capability to integrate many familiar security controls, especially at the network layer.

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 falls onto the provider, but both securing the applications developed against the platform and developing them securely, belong to the consumer.

Figure 6—How Security Gets Integrated

Understanding the impact of these differences between service models and how they are deployed is critical to managing the risk posture of an organization.

IV. CLOUD DEPLOYMENT MODELS

Regardless of the service model utilized (SaaS, PaaS, or IaaS), there are four deployment models for cloud services with derivative variations that address specific requirements.

It is 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 consumer's datacenter, usually via virtual private network (VPN) connectivity.

The architectural mind-set used when designing "solutions" has clear implications on the future flexibility, security, and mobility of the resultant solution, as well as its collaborative capabilities. As a rule of thumb, perimeterized solutions are less effective than de-perimeterized solutions in each of the four areas. Careful consideration should also be given to the choice between proprietary and open solutions for similar reasons.

Deployment models

- ***Public Cloud.*** The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- ***Private Cloud***. The cloud infrastructure is operated solely for a single organization. It may be managed by the organization or by a third party and may be located on-premise or off-premise.
- ***Community Cloud***. The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, or compliance considerations). It may be managed by the organizations or by a third party and may be located on-premise or off-premise.
- ***Hybrid Cloud***. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

V. INFORMATION MANAGEMENT AND DATA SECURITY

The primary goal of information security is to protect the fundamental data that powers our systems and applications. As companies transition to cloud computing, the traditional methods of securing data are challenged by cloud-based architectures. Elasticity, multi-tenancy, new physical and logical architectures, and abstracted controls require new data security strategies. In many cloud deployments, users even transfer data to external — or even public — environments in ways that would have been unthinkable only a few years ago.

Managing information in the era of cloud computing is a daunting challenge that affects all organizations; even those that aren't seemingly actively engaged in cloud-based projects. It begins with managing internal data and cloud migrations and extends to securing information in diffuse,

cross-organization applications and services. Information management and data security in the cloud era demand both new strategies and technical architectures. Fortunately not only do users have the tools and techniques needed, but the cloud transition even creates opportunities to better secure data in our traditional infrastructure.

A. Cloud Information Architectures

Cloud information architectures are as diverse as the cloud architectures themselves. While this section can't possibly cover all potential permutations, there are certain consistent architectures within most cloud services.

Infrastructure as a Service

IaaS, for public or private cloud, generally includes the following storage options:

- **Raw storage.** This includes the physical media where data is stored. May be mapped for direct access in certain private cloud configurations.
- **Volume storage.** This includes volumes attached to IaaS instances, typically as a *virtual hard drive*. Volumes often use *data dispersion* to support resiliency and security.
- **Object storage.** Object storage is sometimes referred to as file storage. Rather than a virtual hard drive, object storage is more like a file share accessed via **API's**[23] or web interface.
- **Content Delivery Network.** Content is stored in object storage, which is then distributed to multiple geographically distributed nodes to improve Internet consumption speeds.

Platform as a Service

PaaS both provides and relies on a very wide range of storage options.

PaaS may provide:

- **Database as a Service.** A multitenant database architecture that is directly consumable as a service. Users consume the database via APIs or direct **SQL**[24] calls, depending on the offering. Each customer's data is segregated and isolated from other tenants. Databases may be relational, flat, or any other common structure.
- **Hadoop/MapReduce/Big Data as a Service.** *Big Data* is data whose large scale, broad distribution, heterogeneity, and currency/timeliness require the use of new technical architectures and analytics. *Hadoop* and other *Big Data* applications may be offered as a cloud platform. Data is typically stored in *Object Storage* or another distributed file system. Data typically needs to be close to the processing environment, and may be moved temporally as needed for processing.
- **Application storage.** Application storage includes any storage options built into a PaaS application platform and consumable via API's that doesn't fall into other storage categories.

PaaS may consume:

- **Databases.** Information and content may be directly stored in the database (as text or binary objects) or as files referenced by the database. The database itself may be a collection of IaaS instances sharing common back-end storage.
- **Object/File Storage.** Files or other data are stored inobject storage, but only accessed via the PaaS API.
- **Volume Storage.** Data may be stored in IaaS volumes attached to instances dedicated to providing the PaaS service.
- **Other.** These are the most common storage models, but this is a dynamic area and other options may be available.

Software as a Service

As with PaaS, SaaS uses a very wide range of storage and consumption models. SaaS storage is always accessed via a web-based user interface or client/server application. If the storage is accessible via API then it's considered PaaS. Many SaaS providers also offer these PaaS APIs.

SaaS may provide:

- **Information Storage and Management.** Data is entered into the system via the web interface and stored within the SaaS application (usually a back-end database). Some SaaS services offer data set upload options, or PaaS API's.
- **Content/File Storage.** File-based content is stored within the SaaS application (e.g., reports, image files, documents) and made accessible via the web-based user interface.

SaaS may consume:

- **Databases.** Like PaaS, a large number of SaaS services rely on database back-ends, even for file storage.

- **Object/File Storage.** Files or other data are stored in object storage, but only accessed via the SaaS application.

- **Volume Storage.** Data may be stored in IaaS volumes attached to instances dedicated to providing the SaaS service.

B. Data (Information) Dispersion

Data (Information) Dispersion is a technique that is commonly used to improve data security, but without the use of encryption mechanisms. These sorts of algorithms (**IDA**[25] for short) are capable of providing high availability and assurance for data stored in the cloud, by means of data fragmentation, and are common in many cloud platforms. In a fragmentation scheme, a file *f* is split into *n* fragments; all of these are signed and distributed to *n* remote servers. The user then can reconstruct *f* by accessing *m* arbitrarily chosen fragments. The fragmentation mechanism can also be used for storing long-lived data in the cloud with high assurance.

When fragmentation is used along with encryption, data security is enhanced: an adversary has to compromise *m* cloud nodes in order to retrieve *m* fragments of the file *f*, and then has to break the encryption mechanism being used.

C. The Data Security Lifecycle

Although *Information Lifecycle Management* is a fairly mature field, it doesn't map well to the needs of security professionals. The Data Security Lifecycle is different from Information Lifecycle Management, reflecting the different needs of the security audience.
The lifecycle includes six phases from creation to destruction. Although it is shown as a linear progression, once created, data can bounce between phases without restriction, and may not pass through all stages (for example, not all data is eventually destroyed).

1. **Create.** Creation is the generation of new digital content, or the alteration/updating/modifying of existing content.

2. **Store.** Storing is the act committing the digital data to some sort of storage repository and typically occurs nearly simultaneously with creation.

3. **Use.** Data is viewed, processed, or otherwise used in some sort of activity, not including modification.

4. **Share.** Information is made accessible to others, such as between users, to customers, and to partners.

5. **Archive.** Data leaves active use and enters long-term storage.

6. **Destroy.** Data is permanently destroyed using physical or digital means (*e.g.*, cryptoshredding).

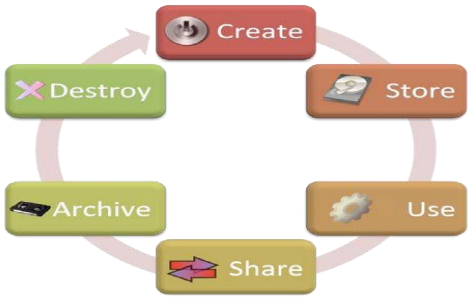

D. Locations and Access

The lifecycle represents the phases information passes through but doesn't address its location or how it is accessed.

Locations

This can be illustrated by thinking of the lifecycle not as a single, linear operation, but as a series of smaller lifecycles running in different operating environments. At nearly any phase data can move into, out of, and between these environments.

Due to all the potential regulatory, contractual, and other jurisdictional issues it is extremely important to understand both the logical and physical locations of data.

Access

When users know where the data lives and how it moves, they need to know who is accessing it and how. There are two factors here: 1. Who accesses the data? 2. How can they access it (device & channel)?

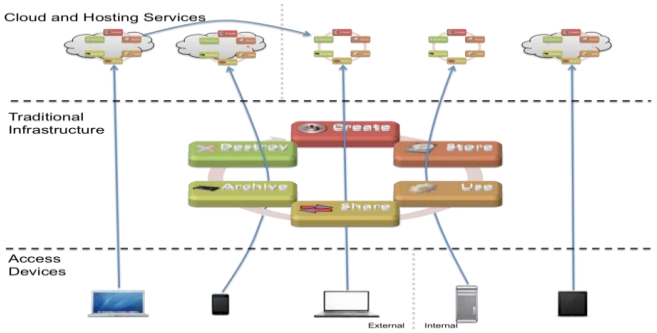

Data today is accessed using a variety of different devices. These devices have different security characteristics and may use different applications or clients.

E. Functions, Actors, and Controls

The next step identifies the *functions* that can be performed with the data, by a given actor (person or system) and a particular location.

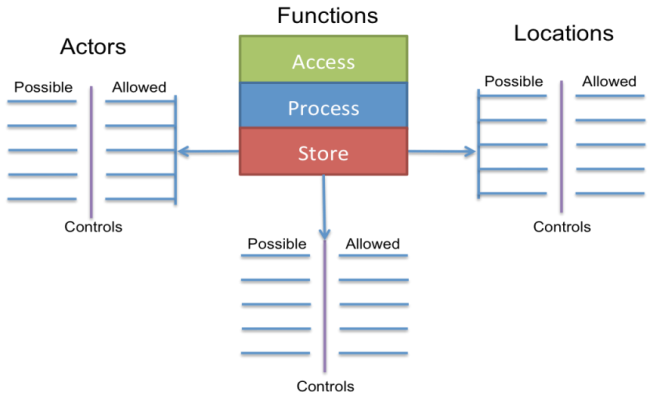

Functions

There are three things we can do with a given datum:

- **Access.** View/access the data, including creating, copying, file transfers, dissemination, and other exchanges of information.
- **Process.** Perform a transaction on the data: update it; use it in a business processing transaction, etc.
- **Store.** Hold the data (in a file, database, etc.).

The table below shows which functions map to which phases of the lifecycle:

Table 1—Information Lifecycle Phases

	Create	Store	Use	Share	Archive	Destroy
Access	X	X	X	X	X	X
Process	X		X			
Store		X			X	

An *actor* (person, application, or system/process, as opposed to the access device) performs each function in a *location*.

Controls

A *control* restricts a list of *possible* actions down to *allowed* actions. The table below shows one way to list the possibilities, which the user then maps to controls.

Table 2—Possible and Allowed Controls

Function		Actor		Location	
Possible	Allowed	Possible	Allowed	Possible	Allowed

VI. INFORMATION GOVERNANCE

Information governance includes the policies and procedures for managing information usage. It includes the following key features:

- **Information Classification.** High-level descriptions of important information categories. Unlike with *data classification* the goal isn't to label every piece of data in the organization, but rather to define high-level categories like "regulated" and "trade secret" to determine which security controls may apply.
- **Information Management Policies.** Policies to define what activities are allowed for different information types.
- **Location and Jurisdictional Polices.** Where data may be geographically located, which also has important legal and regulatory ramifications.
- **Authorizations.** Define which types of employees/users are allowed to access which types of information.
- **Ownership.** Who is ultimately responsible for the information.
- **Custodianship.** Who is responsible for managing the information, at the bequest of the owner.

VII. DATA SECURITY

Data security includes the specific controls and technologies

used to enforce information governance. This has been broken out into three sections to cover detection (and prevention) of data migrating to the cloud, protecting data in transit to the cloud and between different providers/environments, and protecting data once it's within the cloud.

A. Detecting and Preventing Data Migrations to the Cloud

A common challenge organizations face with the cloud is managing data. Many organizations report individuals or business units moving often sensitive data to cloud services without the approval or even notification of IT or security.
Aside from traditional data security controls (like access controls or encryption), there are two other steps to help manage unapproved data moving to cloud services:

1. Monitor for large internal data migrations with Database Activity Monitoring (**DAM**) [26]and File Activity Monitoring (**FAM**)[27].
2. Monitor for data moving to the cloud with URL filters and Data Loss Prevention.

B. Internal Data Migrations

Before data can move to the cloud it needs to be pulled from its existing repository. Database Activity Monitoring can detect when an administrator or other user pulls a large data set or replicates a database, which could indicate a migration.

File Activity Monitoring provides similar protection for file repositories, such as file shares.

C. Movement to the Cloud

A combination of URL filtering (web content security gateways) and Data Loss Prevention (DLP) can detect data moving from the enterprise into the cloud.
URL filtering allows you to monitor (and prevent) users connecting to cloud services. Since the administrative interfaces for these services typically use different addresses than the consumption side, the user can distinguish between someone accessing an administrative console versus a user accessing an application already hosted with the provider.

Look for a tool that offers a cloud services list and keeps it up to date, as opposed to one that requires creating a custom category, and the user managing the destination addresses.

For greater granularity, use Data Loss Prevention. DLP tools look at the actual data/content being transmitted, not just the destination. Thus the user can generate alerts (or block) based on the classification of the data. For example, the user can allow corporate private data to go to an approved cloud service but block the same content from migrating to an unapproved service.

The insertion point of the DLP solution can determine how successfully data leakage can be detected. For example, availability of cloud solutions to various users (e.g., employees, vendors, customers) outside of the corporate network environment avoids or nullifies any DLP solutions if they are inserted at the corporate boundary.

D. Protecting Data Moving To (And Within) the Cloud

In both public and private cloud deployments, and throughout the different service models, it's important to protect data in transit. This includes:

- Data moving from traditional infrastructure to cloud providers, including public/private, internal/external and other permutations.
- Data moving between cloud providers.
- Data moving between instances (or other components) within a given cloud. There are three options (or order of preference):

1. **Client/Application Encryption.** Data is encrypted on the endpoint or server before being sent across the network or is already stored in a suitable encrypted format. This includes local client (agent-based) encryption (e.g., for stored files) or encryption integrated in applications.
2. **Link/Network Encryption.** Standard network encryption techniques including SSL, VPNs, and SSH. Can be hardware or software. End to end is preferable but may not be viable in all architectures.
3. **Proxy-Based Encryption.** Data is transmitted to a proxy appliance or server, which encrypts before sending further on the network. Often a preferred option for integrating into legacy applications but is not generally recommended.

E. Protecting Data in the Cloud

With such a wide range of options and technologies available in cloud computing, there is no way to cover all possible security options. The following are some of the more useful technologies and best practices for securing data within various cloud models.

Content Discovery

Content discovery includes the tools and processes to identify sensitive information in storage. It allows the organization to define policies based on information type, structure, or classification and then scans stored data using advanced content analysis techniques to identify locations and policy violations.

Content discovery is normally a feature of Data Loss Prevention tools; for databases, it is sometimes available in Database Activity Monitoring products. Scanning can be via accessing file shares or a local agent running on an operating system. The tool must be "cloud aware" and capable of working within your cloud environment (e.g., able to scan object storage). Content discovery may also be available as a managed service.

IaaS Encryption

F. Volume Storage Encryption

Volume encryption protects from the following risks:

- Protects volumes from snapshot cloning/exposure
- Protects volumes from being explored by the cloud provider (and private cloud admins)
- Protects volumes from being exposed by physical loss of drives (more for compliance than a real-world security issue)

IaaS volumes can be encrypted using three methods:

- **Instance-managed encryption.** The encryption engine runs within the instance, and the key is stored in the volume but protected by a passphrase or keypair.
- **Externally managed encryption.** The encryption engine runs in the instance, but the keys are managed externally and issued to the instance on request.
- **Proxy encryption.** In this model you connect the volume to a special instance or appliance/software, and then connect your instance to the encryption instance. The proxy handles all crypto operations and may keep keys either onboard or external.

Object Storage Encryption

Object storage encryption protects from many of the same risks as volume storage. Since object storage is more often exposed to public networks, it also allows the user to implement *Virtual Private Storage*. Like a VPN, a **VPS**[28] allows use of a public shared infrastructure while still protecting data, since only those with the encryption keys can read the data even if it is otherwise exposed.

- **File/Folder encryption and Enterprise Digital Rights Management.** Use standard file/folder encryption tools or EDRM to encrypt the data before placing in object storage.
- **Client/Application encryption.** When object storage is used as the back-end for an application (including mobile applications), encrypt the data using an encryption engine embedded in the application or client.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to object storage.

PaaS Encryption

Since PaaS is so diverse, the following list may not cover all potential options:

- **Client/application encryption.** Data is encrypted in the PaaS application or the client accessing the platform.
- **Database encryption.** Data is encrypted in the database using encryption built in and supported by the database platform.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to the platform.
- **Other.** Additional options may include API's built into the platform, external encryption services, and other variations.

SaaS Encryption

SaaS providers may use any of the options previously discussed. It is recommended to use per-customer keys when possible to better enforce multi-tenancy isolation. The following options are for SaaS consumers:

- **Provider-managed encryption.** Data is encrypted in the SaaS application and generally managed by the provider.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to the SaaS application.

Encryption operations should use whatever encryption method is most appropriate, which may include shared keys or public/private keypairs and an extensive **PKI/PKO** [29](Public Key Infrastructure/Operations) structure. Please see

Domain 11 for more information on encryption and key management.

VIII. DATA LOSS PREVENTION

Data Loss Prevention (DLP) is defined as:

Products that, based on central policies, identify, monitor, and protect data at rest, in motion, and in use, through deep content analysis.

DLP can provide options for how data found violation of policy is to be handled. Data can be blocked (stopping a workflow) or allowed to proceed after remediation by encryption using methods such as DRM, ZIP, or OpenPGP.

DLP is typically used for content discovery and to monitor data in motion using the following options:

- **Dedicated appliance/server.** Standard hardware placed at a network chokepoint between the cloud environment and the rest of the network/Internet or within different cloud segments.
- ***Virtual appliance***
- **Endpoint agent**
- **Hypervisor-agent**. The DLP agent is embedded or accessed at the hypervisor level, as opposed to running in the instance.
- **DLP SaaS.** DLP is integrated into a cloud service (e.g., hosted email) or offered as a standalone service (typically content discovery).

IX. Database and File Activity Monitoring

Database Activity Monitoring (DAM) is defined as:

Database Activity Monitors capture and record, at a minimum, all Structured Query Language (SQL) activity in real time or near real time, including database administrator activity, across multiple database platforms; and can generate alerts on policy violations.

DAM supports near real time monitoring of database activity and alerts based on policy violations, such as SQL injection attacks or an administrator replicating the database without approval. DAM tools for cloud environments are typically agent-based connecting to a central collection server (which is typically virtualized). It is used with dedicated database instances for a single customer, although in the future may be available for PaaS.

File Activity Monitoring (FAM) is defined as:

Products that monitor and record all activity within designated file repositories at the user level, and generate alerts on policy violations.

FAM for cloud requires use of an endpoint agent or placing a physical appliance between the cloud storage and the cloud consumers.

Application Security

A large percentage of data exposures are the result of attacks at the application layer, particularly for web applications. Please see Domain 10 for more information on application security.

Privacy Preserving Storage

Almost all cloud-based storage systems require some authentication of participants (cloud user and/or CSP) to establish trust relations, either for only one endpoint of communication or for both. Although cryptographic certificates can offer sufficient security for many of these purposes, they do not typically cater to privacy because they are bound to the identity of a real person (cloud user). Any usage of such a certificate exposes the identity of the holder to the party requesting authentication. There are many scenarios (e.g., storage of Electronic Health Records) where the use of such certificates unnecessarily reveals the identity of their holder.

Over the past 10-15 years, a number of technologies have been developed to build systems in a way that they can be trusted, like normal cryptographic certificates, while at the same time protecting the privacy of their holder (i.e., hiding the real holder's identity). Such attribute-based credentials are issued just like ordinary cryptographic credentials (e.g.,X.509 credentials) using a digital (secret) signature key. However, attribute-based credentials (ABCs) allow their holder to transform them into a new credential that contains only a subset of the attributes contained in the original credential. Still, these transformed credentials can be verified just like ordinary cryptographic credentials (using the public verification key of the issuer) and offer the same strong security.

Digital Rights Management (DRM)

At its core, Digital Rights Management encrypts content, and then applies a series of *rights*. Rights can be as simple as

preventing copying, or as complex as specifying group or user-based restrictions on activities like cutting and pasting, emailing, changing the content, etc. Any application or system that works with DRM protected data must be able to interpret and implement the rights, which typically also means integrating with the key management system.

There are two broad categories of Digital Rights Management:

- **Consumer DRM** is used to protect broadly distributed content like audio, video, and electronic books destined for a mass audience. There are a variety of different technologies and standards, and the emphasis is on one- way distribution.

- **Enterprise DRM** is used to protect the content of an organization internally and with business partners. The emphasis is on more complex rights, policies, and integration within business environments and particularly with the corporate Directory Service.

Enterprise DRM can secure content stored in the cloud well but requires deep infrastructure integration. It's most useful for document based content management and distribution. Consumer DRM offers good protection for distributing content to customers but does not have a good track record with most technologies being cracked at some point.

X. CONCLUSION

The architects and developers involved in providing security to the data and information deployed on the cloud round the clock for both individual and the organizations the process is needed to improve as the opposite players proving a challenge to the service providers and the users of the cloud. But as of now the convenient and successful environment is realized by the user community and hope to continue good in some crucial condition in future too.

REFERENCES

[1] RABIN, M. O. 1989. Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance. J.

[2] ACM, 36(2), 335–348

[3] SECUROSIS. 2011. The Data Security Lifecycle. http://www.securosis.com/blog/data-security-lifecycle-2.0

[5] SECUROSIS. 2011. Understanding and Selecting a Data Loss Prevention Solution. http://www.securosis.com/research/publication/report-data-loss-prevention-whitepaper

[6] SECUROSIS. 2008. Understanding and Selecting a Database Activity Monitoring solution. http://www.securosis.com/research/publication/report-selecting-a-database-activity-monitoring-solution/

[7] CHAUM, D. L. Feb. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. Communications of the ACM, 24 (2), 84-90.

[8] http://msdn.microsoft.com/en-us/library/cc836393.aspx

[9] http://blogs.msdn.com/b/eugeniop/archive/2010/01/12/adfs-wif-on-amazon-ec2.aspx

[10] http://download.microsoft.com/download/6/C/2/6C2DBA25-C4D3-474B-8977-E7D296FBFE71/EC2-Windows%20SSO%20v1%200--Chappell.pdf

[11] http://www.zimbio.com/SC+Magazine/articles/6P3njtcIjmR/Federation+2+0+identity+ecosystem

[12] http://www.datacenterknowledge.com/archives/2009/07/27/cloud-brokers-the-next-big-opportunity/

[13] http://blogs.oracle.com/identity/entry/cloud_computing_identity_and_access

[14] http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

[15] http://www.burtongroup.com

[16] http://www.pkware.com/appnote

[17] http://www.apps.ietf.org/rfc/rfc4880.html