CLOUD COMPUTING

UNIT-I

INTRODUCTION TO CLOUD COMPUTING

In the simplest terms, cloud computing means storing and accessing data and programs over the Internet instead of your computer's hard drive.

What cloud computing is *not* about is your hard drive. When you store data on–or run programs from the hard drive, that's called local storage and computing. Everything you need is physically close to you, which means accessing your data is fast and easy (for that one computer, or others on the local network).

Online social networks and applications

Social network platforms have rapidly changed the way that people communicate and interact. They have enabled the participation in digital communities as well as the representation, documentation and exploration of social relationships.

I personally believe that as "apps" become more sophisticated, it will become easier for users to share their own services, resources and data via social networks. At the same time, social media sites have large number of users all across the globe, and this makes them ideal candidates for cloud adaptation.

Social networks help boost internet usability by storing heavy multimedia content in cloud storage systems. Videos and photographs are the most popular content on social media, which essentially use up the maximum space allocated to them. They have the capacity to slow down applications and servers with all of their resource demands. Cloud computing vendors such as Sales force and Amazon nowadays provide varied services including Customer Relationship Management (CRM) and Enterprise Resource Planning (ERP). As they deliver these things through cloud servers, clients can use the flexibility and scalability of the system without purchasing standalone software or hardware.

Different Clouds

Cloud computing has taken over the Information technology world at a high speed. Cloud computing comes in 3 major forms: private clouds, public clouds and hybrid clouds. Based on the kind of data that one works with, you can compare and use private, public or hybrid clouds with respect to the various levels of management and security needed. Let's look into what each one is:

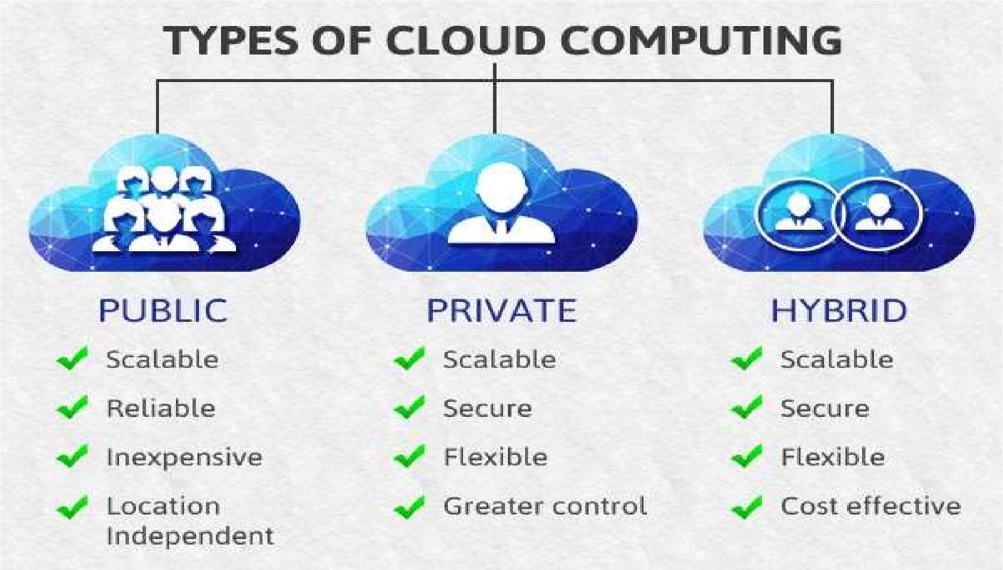

1. **Private Cloud**

   A Private Cloud is implemented using a dedicated data center infrastructure of hardware and software that is used privately by an organization. The data center can be on-premises or off-premises. It is not shared with another organization. If the data center is shared, then its called a Virtual Private Cloud.

   The Cloud Computing Stack in a Private Cloud is dedicated to the organization. If the data center is shared but not the Cloud Computing Stack, that is a Virtual Private Cloud. When both the Cloud Computing Stack and the data centre are shared, then it becomes a Public Cloud. A Private Cloud may participate in a Hybrid Cloud.

2. **Public Cloud**

   A Public Cloud is implemented using a shared data center infrastructure of hardware and software that is shared by multiple organizations. The data center is generally off-premises.
   The Cloud Computing Stack in a Public Cloud is also shared with other organizations. The data, however, for each organization is kept separate. If the data center is shared but not the Cloud Computing Stack, that is a Virtual Private Cloud. When neither the Cloud Computing Stack nor the data center is shared, then that is known as a Private Cloud. A Public Cloud may participate in a Hybrid Cloud.

3. **Hybrid Cloud**

   A Hybrid Cloud is any combination of Clouds. It could be a Private Cloud and one or more Public Clouds. Similarly it could be a Virtual Private Cloud and one or more Public Clouds. But this is so much more than just multiple Clouds. There has got to be resources that are shared among the Clouds. An example of this is Cloud Bursting.

4. **Community Cloud**

   Community Cloud is another type of cloud computing in which the setup of the cloud is shared manually among different organizations that belong to the same community or area. Example of such a community is where organizations/firms are there along with the financial institutions/banks. A multi-tenant setup developed using cloud among different organizations that belong to a particular community or group having similar computing concern.

Cloud Introduction and Overview

History

The concept of sharing resources has been there since 1950’s. It was released by enterprises that purchasing and maintaining computing cost was very expensive. So for the economical reason sharing of resources is important. Here the term comes cloud computing.

Benefits of Cloud Computing

1. Cost: Moving to cloud computing may reduce the cost of managing and maintaining your IT systems. Rather than purchasing expensive systems and equipment for your business, you can reduce your costs by using the resources of your cloud computing service provider.

2**. Security:** Many organizations have security concerns when it comes to adopting a cloud-computing solution. For one thing, a cloud host's full-time job is to carefully monitor security, which is significantly more efficient than a conventional in-house system, where an organization must divide its efforts between a myriad of IT concerns.

3. Speed: Developing in the cloud enables users to get their applications to market quickly.

4. Easily Manageable: Cloud computing offers simplified and enhanced IT maintenance and management capacities by agreements backed by SLA, central resource administration and managed infrastructure. You get to enjoy a basic user interface without any requirement for installation.

5. Performance: Cloud computing services run worldwide over the network of secure data centers which are regularly updated.

Risks

Confidentiality

Probably the main concern, confidentiality is often mentioned as the reason for not embracing cloud computing. If a company's operations require the handling of sensitive data, the protection of these data becomes a priority and a concern. A business might not feel confident in sharing with an external party their vital information. Responsibility for a data leak could be hard to assign when data are handled and transmitted between two parties.

Security

Not just confidentiality, but the entire structure should be evaluated. Where's your data going to be stored? Who will have access to the information? What security measures and protection does the cloud provider offer? Is all information (even when non-sensitive) transmitted in unsecured plaintext or is it encrypted at all times?

Trust

Not all providers are equal. Services through cloud computing may be interrupted by unforeseen events. Outages from a service provider, for example, can happen. Since providers are unable to guarantee no service disruptions will occur, data may not be available 24/7.

Compatibility

Migration to the cloud might pose problems of compatibility with an existing IT infrastructure, or with a company security requirements and organizational policies. Pre-planning is once again crucial in considering all these aspects prior to committing to the change.

CLOUD COMPUTING ARCHITECTURE

Cloud computing architecture refers to the components and subcomponents required for cloud computing. These components typically consist of a front-end platform, back end platforms.

Introduction to cloud computing architecture

The architecture is mainly divides the cloud architecture into two parts:

1. Front End
2. Back End

Each end is connected to others through a network, generally to the Internet. It is the responsibility of the back-end to provide the security of data for cloud users along with the traffic control mechanism.

The cloud technology architecture also consists of front-end platforms called the cloud client which comprises servers, thin & fat client, tablets & mobile devices. The interaction is done through middleware or via web-browser or virtual sessions. According to Jason Bloomberg of ZapThink, the cloud-oriented architecture can essentially be the building block of IoT (Internet of Things) in which anything can be connected to the internet. The cloud architecture is a combination of both services oriented architecture & event-driven architecture. SO cloud architecture encompasses all elements of the cloud environment.

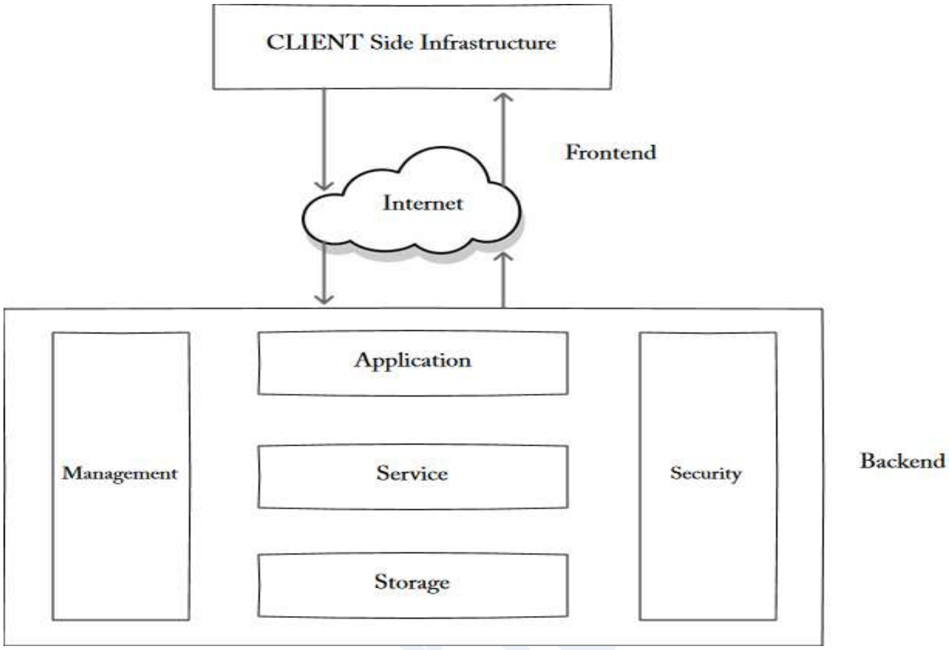

CPU virtualization

CPU virtualization involves a single CPU acting as if it were multiple separate CPUs. The most common reason for doing this is to run multiple different operating systems on one machine. CPU virtualization emphasizes performance and runs directly on the available CPUs whenever possible.

A VM is a duplicate of an existing computer system in which a majority of the VM instructions are executed on the host processor in native mode. Thus, unprivileged instructions of VMs run directly on the host machine for higher efficiency. Other critical instructions should be handled carefully for correctness and stability.

Hypervisor

Hypervisors are a crucial piece of software that makes virtualization possible. Hypervisors are virtual machines that manage multiple operating systems from one piece of physical hardware. These operating systems are referred to as guests, and through the hypervisor's resources, they can be distributed in a variety of ways to meet users' computing needs. For instance, a virtual machine with 4 GB of RAM and 120 GB of hard drive space can easily and instantly be scaled up with the use of hypervisor. A hypervisor is sometimes also called a virtual machine monitor.

Types of Hypervisor:

Type 1 Hypervisor (Also called bare metal or native):

The Type 1 is a layer of software we install directly on top of a physical server and its underlying hardware.

There is no software or any operating system in between, hence the name “bare metal hypervisor.” For this reason, type 1 hypervisors proved to provide excellent performance and stability since they do not run inside Windows or other operating systems.

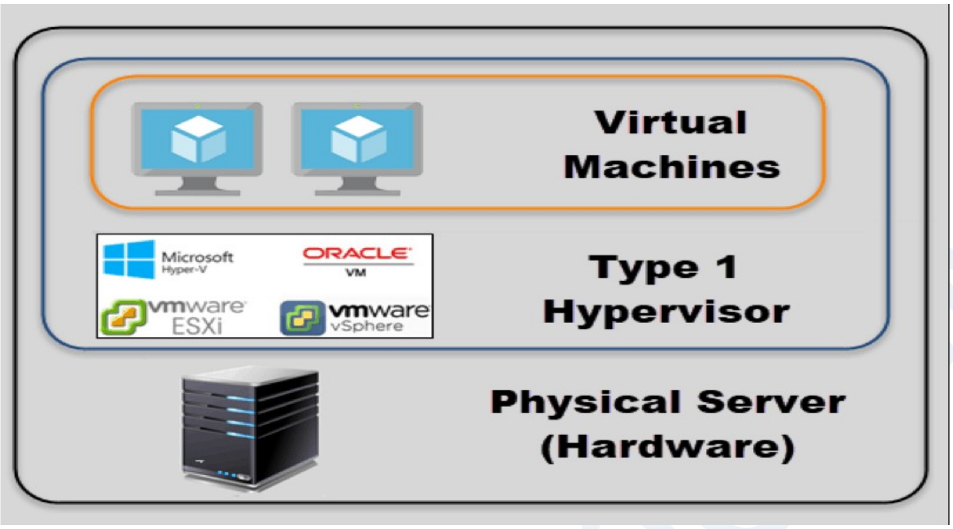

Type 2 Hypervisor (Also known as hosted hypervisors):

This type of hypervisor runs inside of an operating system of a physical host machine. Hosted hypervisors have one software layer underneath.

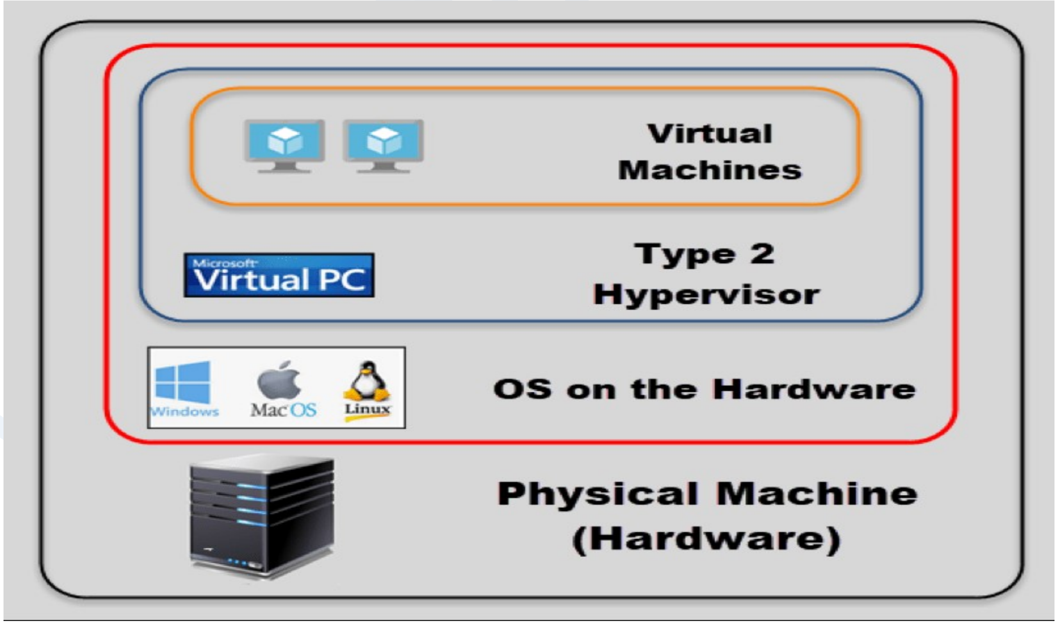

Type 2 hypervisors are typically found in environments with a small number of servers.

e.g. Virtual Box, VMWare etc.

The SPI framework for cloud computing

This acronym stands for the three major services provided through the cloud: software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a service (IaaS).

Software as a Service

In a SaaS model, the customer does not purchase software, but rather rents it for use on a subscription or pay-per-use model. Services provided by this layer can be accessed by end users through Web portals. Therefore, consumers are increasingly shifting from locally installed computer programs to on-line software services that offer the same functionally. This model removes the burden of software maintenance for customers and simplifies development and testing for providers. Example Salesforce.com, which relies on the SaaS model.

Platform as a Service

In a platform-as-a-service (PaaS) model, the service provider offers a development environment to application developers, who develop applications and offer those services through the provider's platform.

A cloud platform offers an environment on which developers create and deploy applications and do not necessarily need to know how many processors or how much memory that applications will be using. Example Google App Engine, an example of Platform as a Service, offers a scalable environment for developing and hosting Web applications, which should be written in specific programming languages such as Python or Java.

Infrastructure as a Service

The IaaS model provides the required infrastructure to run the applications. A cloud infrastructure enables on-demand provisioning of servers running several types of operating systems and a customized software stack. The provider is in complete control of the infrastructure. Infrastructure services are considered to be the bottom layer of cloud computing systems. Example IBM, The definition of infrastructure as a service (IaaS) is pretty simple.

The cloud services delivery model

A cloud delivery model specifies the capabilities offered to users and the applications supported. There are three basic cloud delivery models, Software as a Service (*SaaS*), Platform as a Service (*PaaS*), and Infrastructure as a Service (*IaaS*).

UNIT-II

CLOUD DEPLOYMENT MODELS

Key drivers to adopting the cloud

1. **Increasing business agility –** Business agility is undoubtedly the main benefit and key driver behind Cloud adoption among enterprises. It has been found that enterprises who have adopted Cloud have gained competitive advantage by reducing complexity and increasing business agility. Even for SMBs and independent software vendors (ISVs), business agility has been the main reason for moving to Cloud. Enterprises that have adopted Cloud have observed improved agility due to on-demand self-service and rapid elasticity. Moreover, IT resources can be acquired and deployed more quickly and, once deployed, they can be increased or decreased as needed to meet the demand.
2. **Reducing cost –** This is definitely one of the key drivers behind Cloud adoption. Many enterprises have witnessed a considerable reduction in license and services spend by adopting Cloud services, compared to legacy systems. In many cases it has been found that the cost of replatforming to the cloud is actually much lower than the license renewal of their legacy apps. For SMBs, this benefit leads to a very healthy return on investment (ROI) after year 1. Cloud computing enables organizations to reduce cost through server consolidation, thin clients and community cost sharing.
3. **Enforcing mobility –** An increasing number of enterprises are driven towards Cloud technology, because it is ubiquitous, self-configurable and cost effective. With remote working gaining popularity among organizations, Cloud computing is enabling employees to work at any place, at any time, and on any device. In recent years, the way smart mobile devices market has matured and got acceptance has made mobility a key driver behind Cloud adoption. In the days ahead, more and more enterprises are expected to join the race to launch new Cloud computing solutions of all sizes to be more efficient and gain competitive advantage.
4. **Improving productivity –** Every organization is concerned about improving productivity and Cloud computing is seen as an ideal option by many organizations. Use of Cloud-based tools for email, instant messaging, voice communication, information sharing and development, event scheduling, and conferencing is becoming an increasingly common feature of business life.
5. **Creating new business avenues –** An enterprise can get new business opportunities as a provider of cloud services or added services. Organizations that have good track record of its own IT can become a public Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS) provider. One business opportunity could be, if a company implements a private cloud and has spare capacity, it can sell that additional capacity as public Cloud to another company. On the other hand, software companies can expand their market by providing cloud services in the form of Software-as-a-Service (SaaS).

The impact of cloud computing on users

Cloud computing is changing our lives in many ways. While the technology has been described and commented on at length technically, very few studies have focused on its impact on everyday life.

Social Impact

The likes of YouTube and Google are testimony to a shift in how people are now interacting with others. From remote locations to the global center stage, an event can reach the four corners of the planet by going viral.

Global has reached its true significance, and we’ve seen the emerging of the “citizen journalist” on this global stage. Anyone can turn into an instant reporter, and live news feeds are constantly streaming the media, at times sparking social upheavals.

Education

Educational institutions have been quick to realize the advantages of cloud technology and have been eagerly adopting it for several reasons, including:

- Ability for the students to access data anywhere, anytime, to enrol in online classes and to participate in group activities
- The value of combining business automation processes to streamline subscription, class enrolments and assignment tracking, thus reducing expenses significantly
- Ability for the institutional body to leverage the storage cloud to store the daily 2.5 quintillion bytes of data securely and without the need to cater to a complicated infrastructure
- The benefit of process billing and charging for education and non-education related activities

Development

Cloud technology also offers other benefits to developing countries since they no longer have the burden of investing in costly infrastructures and can tap into data and applications that are readily available in the cloud. I briefly mentioned the education sector above, but the same applies to other areas, such as banking, agriculture, health and science.

Health

There are many reasons why using cloud technology in the healthcare industry is gaining pace. Some examples include: managing non-siloed patient data and sharing it among different parties such as medical professionals or patients checking their own status and treatment follow-ups; reducing operational costs such as data storage; accessing this data through pervasive devices such as mobile phones and going beyond the traditional intranet; implementing a quick solution in a secure environment that is compliant with the Health Insurance Portability and Accountability Act regulations.

Governance in the cloud

Cloud services governance is a general term for applying specific policies or principles to the use of cloud computing services. The goal of cloud services governance is to secure applications and data when they are located remotely.

Cloud governance is a set of rules you create, monitor, and amend as necessary in order to control costs, improve efficiency, and eliminates security risks. There may be other areas of your cloud operations that require governance, but these will become apparent when you first start pulling together the components that will eventually form your rules of governance.

Barriers to cloud computing adoption in the enterprise

Although there are many benefits to adopting cloud computing, there are also some significant barriers to adoption. Two of the most significant barriers to adoption are security and privacy. However, it is important to at least call out what some of the other barriers to adoption are.

Security

Because cloud computing represents a new computing model, there is a great deal of uncertainty about how security at all levels (e.g., network, host, application, and data levels) can be achieved. That uncertainty has consistently led information executives to state that security is their number one concern with cloud computing.

Privacy

The ability of cloud computing to adequately address privacy regulations has been called into question. Organizations today face numerous different requirements attempting to protect the privacy of individuals' information, and it is not clear (i.e., not yet established) whether the cloud computing model provides adequate protection of such information, or whether organizations will be found in violation of regulations because of this new model.

Connectivity and Open Access

The full potential of cloud computing depends on the availability of high-speed access to all. Such connectivity, rather like electricity availability, globally opens the possibility for industry and a new range of consumer products. Connectivity and open access to computing power and information availability through the cloud promote another era of industrialization and the need for more sophisticated consumer products.

Reliability

Enterprise applications are now so critical that they must be reliable and available to support 24/7 operations. In the event of failure or outages, contingency plans must take effect smoothly, and for disastrous or catastrophic failure, recovery plans must begin with minimum disruption. Each aspect of reliability should be carefully considered when engaging with a CSP, negotiated as part of the SLA, and tested in failover drills. Additional costs may be associated with the required levels of reliability; however, the business can do only so much to mitigate risks and the cost of a failure. Establishing a track record of reliability will be a prerequisite for widespread adoption.

SECURITY ISSUES IN CLOUD COMPUTING

Security in cloud computing environment

Cloud computing security refers to the set of procedures, processes and standards designed to provide information security assurance in a cloud computing environment.

Cloud computing security addresses both physical and logical security issues across all the different service models of software, platform and infrastructure. It also addresses how these services are delivered (public, private or hybrid delivery model)

Infrastructure Security:

Infrastructure security is the security provided to protect infrastructure, especially critical infrastructure, such as airports, highways rail transport, hospitals, bridges, transport hubs, network communications, media, the electricity grid, dams, power plants, seaports, oil refineries, and water systems.

The Network Level

When looking at the network level of infrastructure security, it is important to distinguish between public clouds and private clouds. With private clouds, there are no new attacks, vulnerabilities, or changes in risk specific to this topology that information security personnel need to consider. Although your organization's IT architecture may change with the implementation of a private cloud, your current network topology will probably not change significantly. If you have a private extranet in place (e.g., for premium customers or strategic partners), for practical purposes you probably have the network topology for a private cloud in place already. The security considerations you have today apply to a private cloud infrastructure too. And the security tools you have in place (or should have in place) are also necessary for a private cloud and operate in the same way.

However, if you choose to use public cloud services, changing security requirements will require changes to your network topology. You must address how your existing network topology interacts with your cloud provider's network topology. There are four significant risk factors in this use case:

- Ensuring the confidentiality and integrity of your organization's data-in-transit to and from your public cloud provider
- Ensuring proper access control (authentication, authorization, and auditing) to whatever resources you are using at your public cloud provider
- Ensuring the availability of the Internet-facing resources in a public cloud that are being used by your organization, or have been assigned to your organization by your public cloud providers
- Replacing the established model of network zones and tiers with domains

The Host Level

When reviewing host security and assessing risks, you should consider the context of cloud services delivery models (SaaS, PaaS, and IaaS) and deployment models (public, private, and hybrid). Although there are no known new threats to hosts that are specific to cloud computing, some virtualization security threats—such as VM escape, system configuration drift, and insider threats by way of weak access control to the hypervisor—carry into the public cloud computing environment. The dynamic nature (elasticity) of cloud computing can bring new operational challenges from a security management perspective. The operational model motivates rapid provisioning and fleeting instances of VMs. Managing vulnerabilities and patches is therefore much harder than just running a scan, as the rate of change is much higher than in a traditional data center.

The Application Level

Studies indicate that most websites are secured at the network level while there may be security loopholes at the application level which may allow information access to unauthorized users. Software and hardware resources can be used to provide security to applications. In this way, attackers will not be able to get control over these applications and change them. XSS attacks, Cookie Poisoning, Hidden field manipulation, SQL injection attacks, DoS attacks, and Google Hacking are some examples of threats to application level security which resulting from the unauthorized usage of the applications.

Data Security

Majority of cloud service providers store customers' data on large data centres. Although cloud service providers say that data stored is secure and safe in the cloud, customers' data may be damaged during transition operations from or to the cloud storage provider. In fact, when multiple clients use cloud storage or when multiple devices are synchronized by one user, data corruption may happen. Cachin and his colleagues (2009) proposed a solution, Byzantine Protocols, to avoid data corruption. In cloud computing, any faults in software or hardware that usually relate to inappropriate behaviour and intrusion tolerance are called Byzantine fault tolerance (BFT). Scholars use BFT replication to store data on several cloud servers, so if one of the cloud providers is damaged, they are still able to retrieve data correctly. In addition, different encryption techniques like public and private key encryption for data security can be used to control access to data. Service availability is also an important issue in cloud services. Some cloud providers such as Amazon mentions in their licensing agreement that it is possible that their service is not available from time to time. Backups or use of multiple providers can help companies to protect services from such failure and ensure data integrity in cloud storage.

Aspects of Data Security

Data Security Mitigation

If prospective customers of cloud computing services expect that data security will serve as compensating controls for possibly weakened infrastructure security, since part of a customer's infrastructure security moves beyond its control and a provider's infrastructure security may (for many enterprises) or may not (for small to medium-size businesses, or SMBs) be less robust than expectations, you will be disappointed. Although data-in-transit can and should be encrypted, any use of that data in the cloud, beyond simple storage, requires that it be decrypted. Therefore, it is almost certain that in the cloud, data will be unencrypted. And if you are using a PaaS-based application or SaaS, customer-unencrypted data will also almost certainly be hosted in a multitenancy environment (in public clouds).

So, what should you do to mitigate these risks to data security? The only viable option for mitigation is to ensure that any sensitive or regulated data is not placed into a public cloud (or that you encrypt data placed into the cloud for simple storage only). Given the economic considerations of cloud computing today, as well as the present limits of cryptography, CSPs are not offering robust enough controls around data security.

Provider Data and Its Security

In addition to the security of your own customer data, customers should also be concerned about what data the provider collects and how the CSP protects that data. Specifically with regard to your customer data, what metadata does the provider have about your data, how is it secured, and what access do you, the customer, have to that metadata? As your volume of data with a particular provider increases, so does the value of that metadata.

Additionally, your provider collects and must protect a huge amount of security-related data. For example, at the network level, your provider should be collecting, monitoring, and protecting firewall, intrusion prevention system (IPS), security incident and event management (SIEM), and router flow data. At the host level your provider should be collecting system logfiles, and at the application level SaaS providers should be collecting application log data, including authentication and authorization information.

UNIT-III

IDENTITY AND ACCESS MANAGEMENT

According to Gartner, Identity and Access Management (IAM) is the security discipline that enables the right individuals to access the right resources at the right times for the right reasons. IAM addresses the mission-critical need to ensure appropriate access to resources across increasingly heterogeneous technology environments.

Trust boundaries and IAM

In a typical organization where applications are deployed within the organization's perimeter the "trust boundary" is mostly static and is monitored and controlled by the IT department. In that traditional model, the trust boundary encompasses the network, systems, and applications hosted in a private data centre managed by the IT department (sometimes third-party providers under IT supervision). And access to the network, systems, and applications is secured via network security controls including virtual private networks (VPNs), intrusion detection systems (IDSs), intrusion prevention systems (IPSs), and multifactor authentication.

To compensate for the loss of network control and to strengthen risk assurance, organizations will be forced to rely on other higher-level software controls, such as application security and user access controls. These controls manifest as strong authentication, authorization based on role or claims, trusted sources with accurate attributes, identity federation, single sign-on (SSO), user activity monitoring, and auditing. In particular, organizations need to pay attention to the identity federation architecture and processes, as they can strengthen the controls and trust between organizations and cloud service providers (CSPs).

Challenges faced by IAM

New cloud-based identity and access management (IAM) services are growing in popularity as more organizations are opting for them to provide a unified and simple identity management. They may add extra security and protection to your company resources. But, it poses key challenges like proper assessment of the existing IT infrastructure.

1. Identity Provisioning / De-provisioning

This concerns with providing a secure and timely management of on-boarding (provisioning) and off-boarding (de-provisioning) of users in the cloud.

When a user has successfully authenticated to the cloud, a portion of the system resources in terms of CPU cycles, memory, storage and network bandwidth is allocated. Depending on the capacity identified for the system, these resources are made available on the system even if no users have been logged on.

2. Maintaining a single ID across multiple platforms and organizations

It is tough for the organizations to keep track of the various logins and ID that the employees maintain throughout their tenure. The centralised federated identity management (FIdM) is the answer for this issue. Here users of cloud services are authenticated using a company chosen identity provider (IdP).

Compliance Visibility: Who has access to what

When it comes to cloud services, it's important to know who has access to applications and data, where they are accessing it, and what they are doing with it. Your IAM should be able to provide a centralised compliance reports across access rights, provisioning/de-provisioning, and end-user and administrator activity. There should be a central visibility and control across all your systems for auditing purposes.

Security when using 3rd party or vendor network

A lot of services and applications used in the cloud are from 3rd party or vendor networks. You may have secured your network, but can't guarantee that their security is adequate.

Relevant IAM standards and protocols for cloud services

There are several standards and protocols which can be used for the identity management in the cloud environments. The various identity management protocols differ in their features such as the data format supported, protocols to exchange credentials between the entities involved etc. some of the well-known identity management protocols which help in establishing a federation among the individual partners are given below:

1. **Security Assertion Markup Language (SAML):** The SAML includes a set of specifications for exchanging the authentication, authorization and attribute assertions across the federation. SAML uses XML-based data format and this protocol is managed under the Organization for the Advancement of Structured Information Standards (OASIS). The various use cases of SAML include SSO in the federation, identity/account linkage, session management and secure web services.
2. **Service Provisioning Markup Language (SPML):** This is an XML-based security framework developed by OASIS and it could be used to exchange information related to the users, resources and service provisioning within a group of cooperative organizations. This helps in automating the provisioning and de-provisioning of user accounts with the CSP.
3. **eXtensible Access Control Markup Language (XACML):** in order to implement an access control mechanism, this XML-based access control language developed by OASIS could be used. This language provides the XML-schema which could be used to protect the resources by making access decisions over these resources.
4. **OpenID:** OpenID provides a user-centric identity framework for the authentication purposes. OpenID 2.0 supports the identification of users through URL or XRI addresses. OpenID uses the concept of Relying Party (RP) and OpenID providers. User's authentication data are stored by the OpenID provider and user has the flexibility to decide who have access to his authentication data maintained by the OpenID provider.
5. **OAuth:** OAuth is an open source Identity Management protocol that could be used to provide the authorization of users across different applications without disclosing the user's identity credentials. Identity tokens issued by the identity provider are used by the third-party applications to gain access to the user's protected data.
6. **WS-Federation:** WS-Federation is a part of the web services security specifications and is meant for the federation of applications or web services. WS-Federation specifications are extensions to WS-Trust protocol. This protocol can be used to share the identity information of various users across multiple security domains and organizations.

Cloud Authorization Management

Medium-size and large organizations usually have specific requirements for authorization features for their cloud users. In some cases, a business application may require role-based access control (RBAC), in which case authorization is structured to suit the organization’s functional role requirements. As of this writing, cloud service authorization enforcement and management capabilities are weak and when they are available, they are very coarse-grained. The services available may not meet your enterprise requirements.

Most cloud services support at least dual roles: administrator and end user. It is a normal practice among CSPs to provision the administrator role with administrative privileges. These privileges allow administrators to provision and deprovision identities, basic attribute profile and in some cases to set access control policies such as password strength and trusted networks from which connections are accepted.

IAM Support for Compliance Management

As much as cloud IAM architecture and practices impact the efficiency of internal IT processes, they also play a major role in managing compliance within the enterprise. Properly implemented IAM practices and processes can help improve the effectiveness of the controls identified by compliance frameworks. E.g. by automating the timely provisioning and deprovisioning of users and entitlements, organizations can reduce the risk of unauthorized users accessing cloud services and meet your privacy and compliance requirements. In addition, identity and attribute management will be key areas of compliance focus for regulatory and privacy issues-proper IAM governance processes should be instituted to address these issues.