CLOUD COMPUTING
QP CODE: 559
SUBJECT CODE: 35271

PART –A

(Each question carries 2 marks, Answer any FIVE questions, Q.No. 8 – Compulsory)

1. What is cloud computing?

Cloud computing is the process of delivering or providing computational resources like software and/or hardware as a service over the cloud (internet). It is also known as internet computing.

2. What is autonomic computing?

Autonomic computing refers to the self managing characteristics of distributed computing resources, adapting to unpredictable changes. It controls the functioning computer applications and systems without input from the user. This computing model has systems that run themselves, capable of doing high level functions.

3. What is SPI?

The acronym for SPI stands for three major services provided through the cloud.
They are as follows,

1. Software as a Service (SaaS)
2. Platform as a Service (PaaS)
3. Infrastructure as a Service (IaaS)

4. State any two service provider of SaaS.

Some of the service providers are

1. Amazon Web services
2. Google Apps
3. icloud
4. Oracle
5. Salesforce.com
6. Windows Azure

5. What is virtualization?

Virtualization is the key component of cloud computing for providing computing and storage services. Virtualization is the ability to run multiple operating systems on a single physical system and share the underlying hardware resources. It is the process by which one computer hosts the appearance of many computers.

6. What is thin client?

This refers to computers that do not have internal hard drives. They allow the server do all the work, but then display the information on the screen.

7. What is storage networking?

Storage networking is the practice of linking together storage devices and connecting them to other IT networks. Storage networks provide a centralized repository for digital data that can be accessed by many users, and they use high-speed connections to provide fast performance.

The phrase "storage networking" is commonly used in reference to storage area networks (SANs).

8. What is CSA?

CSA is an industry working group that studies security issues in cloud computing and offers recommendations to its members. It gives the cloud computing stack model, which shows how different functional units in a network stack relate to one another.

PART B

(Each question carries 3 marks, Answer any FIVE questions. Q.No 16 – Compulsory)

9. State any three essential characteristics of cloud computing.

On-Demand self service

It is one of the essential characteristic of cloud that allows user to receive the services such as computing resources, server time and network storage automatically without direct interaction with the service provider. The applications and resources can be assigned and removed within minutes using cloud catalogs. Some of the popular on demand self service providers are AWS (Amazon Web Services), Google, Microsoft, IBM, Salseforce.com.

Broad network access

This is another essential aspect that is available over the network. They are accessed by using standard mechanisms in thick or thin client platforms.

Location independent resource pooling

The service provider's resources are pooled in order to serve multiple consumers. There is a sense of location independence as the customer has no control over location where the resources are provided. Consumers need not worry about how the cloud allocates the provided resources.

Rapid elasticity

The definition of elasticity is the ability to scale the resources up and down as required. The storage on cloud seems to be unlimited for the client. The consumer can use as much as he needs at any time.

Measured services

Another essential attribute is that the resources can be measured, controlled and reported. This provides transparency for both provider and consumer of the used service. Metering capability is used to control and optimize resource use.

10. State the benefits of cloud computing.

Scalability:

The ability of a model to be extended to manage the amount of work growth in an effective manner is called scalability. Cloud-computing resources can be rapidly scaled according to subscribers convenience. If there is a sudden necessity for more computer resources, instead of buying new equipment we can buy additional resources form cloud providers. After the endeavor is over we can stop using those services.

Simplicity:

In most cases cloud-computing is free to use. It is very simple that users can easily understand which is the biggest advantage of cloud-computing. It is possible to get our application started instantly.

Vendors:

The service providers are called vendors. Some of the well known vendors are Google, Amazon, Microsoft, IBM. These providers offer reliable services to their customers.

Security:

There are also some risks when using a cloud vendor. But the reputed firms work hard to keep their consumers data safe and secure. They use complex cryptographic algorithms to authenticate users. To make it even more secure we can encrypt our information before storing it in cloud.

11. What is PaaS?

PaaS is another cloud delivery model which delivers more than just infrastructure. It provides solution stack which is an integrated set of software. This model provides everything a developer needs to build an application for software development and run time.

12. What is Google app engine?

Google app engine is a SaaS provider which was introduced in 2008. It was quite unique cloud system compared to other systems. It provides platform to create applications. It provides infrastructure for hosting. Many high level services which needs to be build are available when using an App Engine.

13. State the types of hardware virtualization?

Different types of hardware virtualization are :

1. Full virtualization
2. Partial virtualization

3. Para virtualization

14. State the limitations of virtualization.

1. If the CPU does not allow for hardware virtualization we can run some operating system in software virtualization but it is generally slower. Some operating system will not run in software virtualization and require to have CPU with hardware virtualization so it would cost more if CPU with hardware virtualization is not possible.

2. If we want a own server and intend to resell a virtual server then it cost high. This mean purchase of 64 bit hardware with multiple CPU's and multiple hard drives.

3. Some of the limitations are in analysis and planning which problems can be divided into three types they are

a. Technical limitation

b. Marketing strategies

c. Political strategies

4. It has a high risk in physical fault.

5. It is more complicated to set up and manage virtual environment with high critical servers in a production environment. It is not easy as managing physical servers.

6. It does not support all applications.

15. What is ISCSI?

Internet Small Computer Systems Interface is an Internet Protocol (IP)-based storage networking standard for linking data storage facilities. It provides block-level access to storage devices by carrying SCSI commands over a TCP/IP network. ISCSI is used to facilitate data transfers over intranets and to manage storage over long distances. It can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet and can enable location-independent data storage and retrieval.

16. What is the goal of encrypted cloud storage?

The goal of encrypted cloud storage is to create a virtual private storage system that maintains confidentiality and data integrity. Encryption should separate stored data (data at rest) from data in transit. Depending upon the particular cloud provider, such as Microsoft allows up to five security accounts per client, and can use these different accounts to create different zones. On Amazon Web Service, we can create multiple keys and rotate those keys during different sessions.

PART-C

(Each question carries marks, Answer division (a) or division (b))

17. (a) (i) Write short notes on origins of cloud computing.

- This concept of providing resources via global networks starts from1960's.
- This idea was introduced by J.C.R.Licklider, who was responsible for ARPANET in 1969.
- Since then the cloud computing has developed a lot. After the wide internet usage in 1990's, large bandwidth was offered. Thus this concept came to use by public.
- One of the first milestones for cloud-computing was the introduction of salesforce.com in 1999. It introduced the concept of delivering enterprise applications through a simple website. This helped software companies to deliver applications over the cloud.
- Amazon Web Services in 2002 was the next major development in cloud computing, which provided a group of cloud based services.
- Amazon launched its Elastic Compute Cloud (EC2) in 2002 as a commercial web service which allows individuals and small scale industries to hire/rent computers on which they can run their own applications.
- Another big milestone was the introduction of web2.0 in 2009. Web2.0 is a website that allows the users to interact with each other in a social media creator. Examples of Web2.0 include social networking websites, blogs etc.
- Google and others started to offer browser based enterprise applications, through services such as Google apps.

- Other important factors that have enabled cloud computing to evolve are virtualization technology, development of universal high-speed bandwidth and universal software standards.
- Most of the IT professionals use cloud computing as it offers increased storage space, high flexibility and very low cost.Thus cloud computing has brought enormous benefits for users.

(ii) Explain briefly the security concerns of cloud computing.

In cloud computing world, security is a two sided coin. The security is very important particularly when moving critical applications and sensitive data to public and shared environments.

Privacy concern with a third party

The important security concern is for privacy considerations. That is, if third party is hosting all our data, we do not know if it is safe or not. Everything that is placed on cloud can be accessed by anyone. There are also other privacy concerns because government can get the data that is placed on cloud easily from organization's servers. Though there are popular companies who provide good security to keep the data safe, it can be hacked. The best procedure is not to perform critical tasks on a cloud platform without extensive security. If it cannot be managed then it is advisable to have less critical data on cloud.

Security level of third party

Service providers are doing all they can to protect their customer's data. As a matter of fact, the vendors will have to make sure that the subscriber has been fully satisfied from their service or else the firm will not be gaining customers. Most of the security problems are due to loss of control, lack of trust and multi-tenancy. Multi-tenancy – it is an architecture in which single instance of a software application serves multiple customers. Each customer is called tenant. These problems exist mainly in third party management models. So there should be strong protection measures in order to prevent the hacking of data.

Security Benefits

Providers do endeavor to ensure security. Cloud provide some of the security measures ensuring the customers data are safe:

Centralised Data

There are some good security traits that come with centralizing your data, making your system more inherently secure.

Reduced Data Leakage:

If the data is centralized and the various devices used like laptop, notebook computers can access the data, no need to backup the data. There is threat for theft of the handheld devices. If the data are lost and although any security measures like encryption is applied and it may be compromised and the entire data may be in the hands of the thief. Moreover by maintaining data on the cloud, employing strong access control, limiting the employee downloading to only what they need to perform a task, computing can limit the amount of information that could be potentially be lost.

Monitoring benefits:

Central storage is easier to control and monitor. The flipside is the nightmare scenario of comprehensive data theft. If your data is maintained on a cloud, it is easier to monitor security than have to worry about the security of numerous servers and clients. The security professional figuring out smart ways to protect and monitor access to data stored in one place (with the benefit of situational advantage) than trying to figure out all the places where the company data resides. You can get the benefits of Thin Clients today but Cloud Storage provides a way to centralize the data faster and potentially cheaper. The logistical challenge today is getting Terabytes of data to the Cloud in the first place.

Instant Swap over - if a server in the Cloud gets compromised (i.e. broken into), then clone that server at the click of a mouse and make the cloned disks instantly available to the Cloud Forensics server. When the swap over is performed its seamless to the users. No need to spend time to replicate the data or fix the breach. Abstracting the hardware allows to do it instantly.

Logging

In cloud logging is improved. Logging is often an afterthought, to solve the issues insufficient disk space is allocated. Cloud Storage changes all this - no more 'guessing' how much storage you need for standard logs. With your logs in the Cloud you can leverage Cloud Compute to index those logs in realtime and get the benefit of instant search results. This help to Compute instances and to measure in and scale as needed based on the logging load - meaning a true real-time view. Most modern operating systems offer extended logging in the form of a C2 audit trail. This is rarely enabled for fear of performance degradation and log size. Now you can 'opt-in' easily - if you are willing to pay for the enhanced logging, you can do so. Granular logging makes compliance and investigations easier.

Secure builds

When you developed your own network and you have to buy third-party security software to get the level of protection you want. With the cloud solution, those tools can be bundled in and available to you and you can develop your system with whatever level of security you desire.

Easier to test impact of security changes: this is a big one. Spin up a copy of your production environment, implement a security change and test the impact at low cost, with minimal startup time. This is a big deal and removes a major barrier to 'doing' security in production environments.

Drive vendors to create more efficient security software:

Billable CPU cycles get noticed. More attention will be paid to inefficient processes; e.g. poorly tuned security agents. Process accounting will make a Comeback as customers target 'expensive' processes. Security vendors that understand how to squeeze the most performance from their software will win.

Security Testing

Reduce cost of testing security**:** A SaaS provider only passes on a portion of their security testing costs. It is shared among the cloud users. The end results is that because you are in a pool with others but you never see the other users but you realize the lower cost for testing. Even with Platform as a Service (PaaS) where your developers get to write code, but the cloud code –scanning tools check for security weakness.

(OR)

17. (b) Discuss the regulatory issues of cloud computing and the government policies.

Regulatory Issues

In the case of cloud computing, regulation might be exactly what we need. Without some rules in place there are chances for unsecure with service or even shifty enough to make off with your data.

'Sensitive Data' is defined as personal information that relates to:

(a)passwords;

(b)financial information such as Bank account or credit card or debit card or other payment instrument details;

(c)physical, psychological and mental health condition;

(d)sexual orientation;

(e)medical records and history;

(f)biometric information;

any detail relating to the above received by the body corporate for provision of services; or any information relating to (a) – (g) that is received, stored or processed by the body corporate under a lawful contract or otherwise. No existing regulation: currently there is no existing regulation. While comparing cloud service providers to banks there are similarities.

Banks deal with money whereas cloud service providers deal with data, both are immense value to consumers and organizations alike.

Location of Stored Data – Service providers generally do not disclose the location where the service subscriber's data are stored. It leaves users in the dark regarding the extent of protection applied to their critical information. Although security certifications could lessen the user's anxiety, the matter of determining if the provider's compliance with legal and regulatory laws includes those that cover the geographical location where data is stored, aside from the laws of the areas where the data was collected.

If government can figure out a way to safeguard data, either from loss or theft of any company facing such a loss would applaud the regulation. One such example is the greatest bank failure in American History. In 2008 the United States government took control of Washington Mutual. On the other hand, there are those who think the government should stay out of it and let competition and market forces guide cloud computing.

Law enforcement agencies have easier access to personal information on cloud data than that stored on a personal computer. Also the big problem is that people using cloud services are not aware of the privacy and security implication on their online email accounts, their LinkedIn account, their MySpace page, and so forth. While these are popular sites for individuals, they are still considered cloud services and their regulation may affect other cloud services.

Government Procurement

There are also questions about whether government agencies will store their data on the cloud. Procurement regulations will have to change for government agencies to be keen on jumping on the cloud. The General Service Administration (GSA) is making a push toward cloud computing, in an effort to reduce the amount of energy their computers consume. The GSA is working with a vendor to develop an application that will calculate how much energy government agencies consume.

Government Policies:

The aim of the cloud policy of government is to realise a comprehensive vision of a government cloud (GI Cloud) environment available for use by central and state government line departments, districts and municipalities to accelerate their ICT-enabled service improvements. As per the guidelines, both cloud service provider (CSP) and government department will have to share responsibility for the managing services provisioned using cloud computing facility.

To implement the policy, Government of India has made an initial step "GI Cloud" which has been coined as 'Meghraj'. The focus of this initiative is to accelerate delivery of e-services in the country while optimizing the expenditure of the Government. The ministry of electronics and IT has issued on important guideline regarding the location of data as follows:

"The terms and conditions of the Empanelment of the Cloud Service Provider has taken care of this requirement by stating that all services including data will be guaranteed to reside in India". The cloud computing service enables its user to hire or use software, storage, servers as per requirement instead of purchasing the whole system. Meity(Ministry of Electronics and Information Technology) has empanelled the following companies for providing cloud computing services to government departments :

1. Microsoft Corp.,
2. Hewlett Packard,
3. IBM India ,
4. Tata Communications,
5. Bharat Sanchar Nigam Ltd (BSNL),
6. Net Magic IT Services,
7. Sify Technologies and
8. CtrlS Data Centers.

The architectural vision of GI Cloud as mentioned above consists of a set of discrete cloud computing environments spread across multiple locations, built on existing or new (augmented) infrastructure as given below :

Components of Meghraj

1. Setting up of State and National Clouds
2. Set up an e-Gov Appstore
3. Empanelment of Cloud Service Providers
4. Empanelment of Cloud Auditors
5. Setting up of Cloud Management Office
    - Setting up an eco-system for Cloud proliferation (Policies, Guidelines, templates, security norms, certification, business models for applications, tariff & revenue models for private sector Cloud services)
    - Awareness workshops, training programs and migration support for cloud adoption by departments.
6. MeghRaj (GI-Cloud) service Directory
7. Setting up of Clouds by other Government entities

Cloud Deployment Models:

The empanelment of the Cloud service offerings of CSPs has been done for a combination of the Cloud Deployment models and Service models as mentioned below:

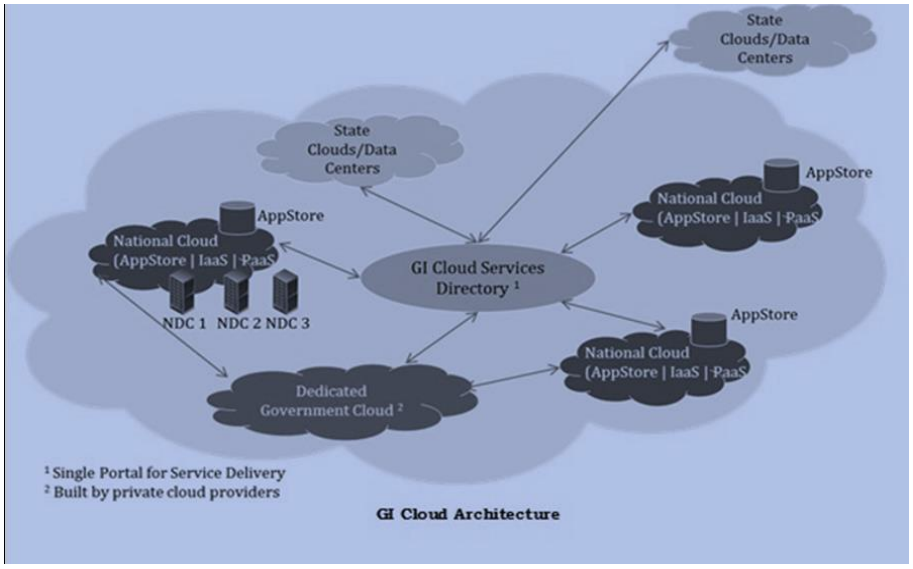

GI Cloud Architecture

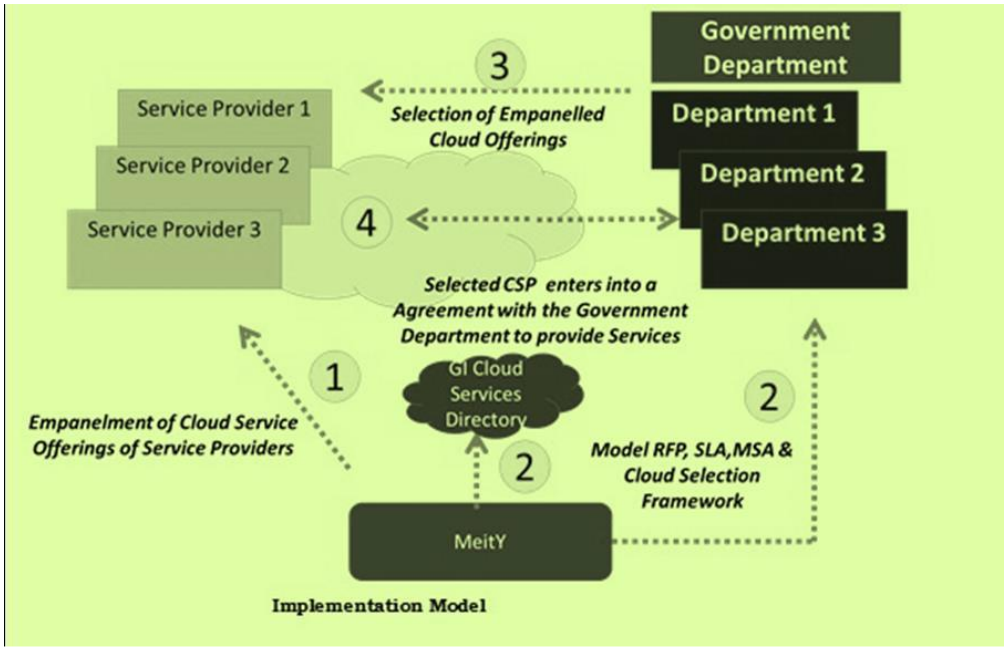

Implementation model

Cloud Deployment Models:

The empanelment of the Cloud service offerings of CSPs has been done for a combination of the Cloud Deployment models and Service models as mentioned below:

Public Cloud :

A shared multi-tenant IT infrastructure is made available over the internet. It is owned and operated by a Cloud Service Provider delivering cloud services to the Government Department.

Virtual Private Cloud :

A logically separated Cloud Infrastructure (Servers, Storage, Network infrastructure and Networks) to protect data, applications and servers and provide robust virtual isolation for the Government Department.

Government Community Cloud

A cloud with IT infrastructure resources which will be dedicated for two or more Government Departments that have common privacy, security and regulatory considerations.

18. (a) Explain in detail about cloud delivery model.

Cloud computing services can be delivered to customers(users) in different ways and depicted in Fig . The cloud computing delivery models include infrastructure, platform and software. These services are provided and used over the internet.

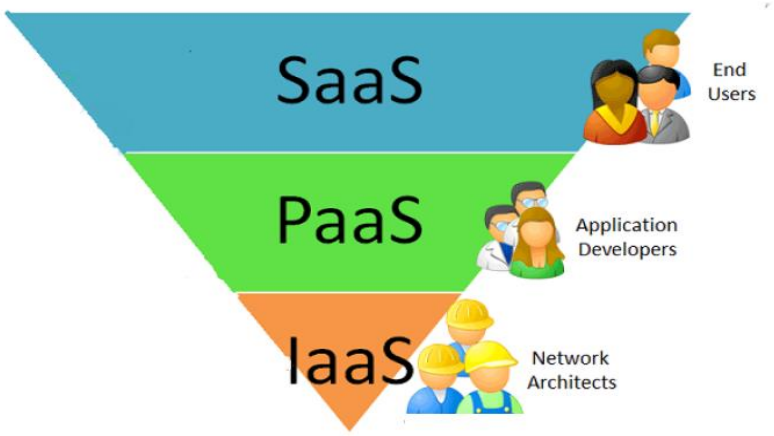

Cloud Delivery Model

Infrastructure as a Service (IaaS)

It is one of the cloud delivery model which provides computer infrastructure or hardware like servers, networking technology, storage as a service. It may also include the delivery of operating system and virtualization technologies to manage these resources.

Platform as a Service (PaaS)

PaaS is another cloud delivery model which delivers more than just infrastructure. It provides solution stack which is an integrated set of software. This model provides everything a developer needs to build an application form software development and run time.

Software as a Service (SaaS)

SaaS is a delivery model which delivers business application designed for a specific purpose. This service is provided over the internet which eliminates the need to install and run the applications on consumer's own computers. It simplifies maintenance and support.

Some of the properties of SaaS are:

- Network or online access
- Centralized management
- Powerful communication features

SaaS works on two distinct modes, they are

- Single multi tenancy
- Fine-grain multi tenancy

(OR)

18. (b) Discuss the operational and economic benefits of SaaS.

Operational benefits of SaaS :

SaaS can improve the consumer's organization effectiveness based on the following benefits,

1. **Managing business driven IT project:**

A SaaS model provides the necessary infrastructure and thus leads to technology projects that address true business needs.

2. **Increasing consumer demand :**

SaaS model provides reliability to deliver near perfect 99.99% system availability. So any number of users can access the system at anytime from anywhere.

3. **Addressing growth :**

This model provides scalability that is easily supported by an increasing number of consumers to meet their own objectives.

4. **Serving new markets quickly and easily:**

SaaS allows the organization to quickly and easily adds programs so as to adapt the changes based on the demand at a faster rate.

5. **On demand:** The solution is self serve and available for use as needed.

6. **Scalable:** It allows for the infinite scalability and quick processing time.

Economic benefits of SaaS :

SaaS not only saves time but also has greater financial benefits.

1. It reduces IT expenses.
2. The implementation cost of SaaS is much lower than the traditional software.
3. It redirects savings expenses towards business improvements.
4. It strengthens the financial capability.
5. By utilizing SaaS, we are free to use as much of any software as we need. This gives you easy and economical access to many programs.
6. SaaS vendors release upgrades for their software, thus users need not put any effort into installing and upgrading the software.
7. Another main benefit in SaaS is that it can quickly and easily be accessed from anywhere by using a web browser.

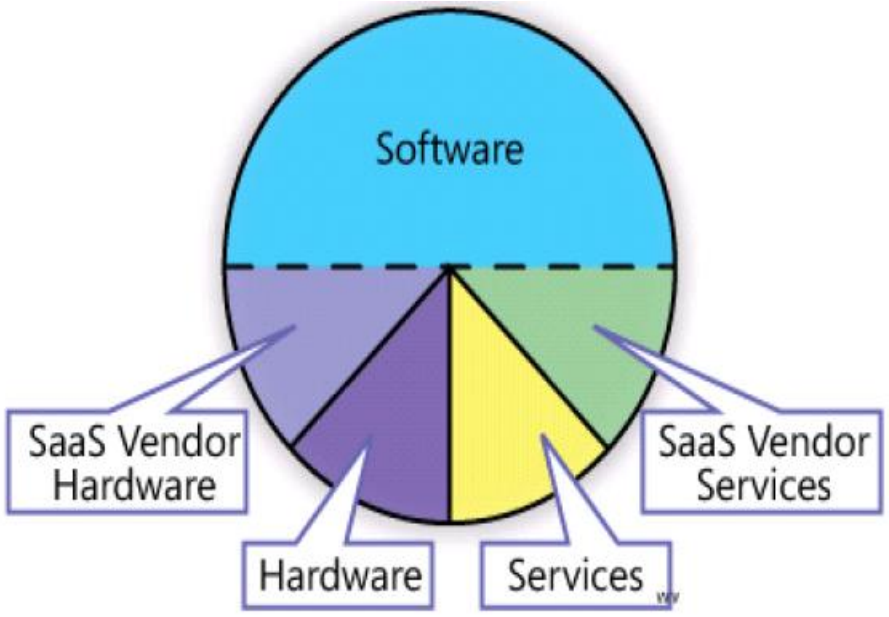

SaaS utilization

19. (a) Explain in detail the various aspects for the need of virtualization in cloud computing.

Virtualization is needed for the following reasons

Costs

With virtualization, administration becomes a lot easier, faster and cost effective. Virtualization lowers the existing cost. It dramatically simplifies the ownership and administration of their existing IT servers. The operational overhead of staffing, backup, hardware and software maintenance has become verify significant in IT budgets and business. In such case virtualization reduces these costs beneficially. By using virtualization we can save the operational costs. Virtualization concentrates on increasing utilization and consolidation of equipment thereby reducing capital costs, cabling, operational costs such as power, cooling, maintenance cost of hardware and software.

Thus Virtualization is cost effective.

Administration

Administrating virtualization has to be done in efficient manner since all the resources are centralized security issues has to be categorized more sensitively. The users access the resources like data storage, hardware or software has to be allocated properly. Since more users will utilize the resources, the sharing of needed resources is complicated.

Administration of virtual server is done through virtual server administration website. By using this virtual server is assigned to application for access.

In this, virtual IP addresses are configured on the load balancer. When a request is sent from the user from certain port on a virtual IP load balancer, it distribute the incoming request among multiple server the needed service will be provided to particular user.

Fast Deployment

Deployment of consolidated virtual servers, migrating physical, servers has to be done.

Virtualization deployment involves several phases and planning. Both server and client systems can support several operating systems simultaneously, virtualization providers offer reliable and easily manageable platform to large companies.

It can be built with independent, isolated units which work together without being tied to physical equipment. Virtualization provides much faster and efficient way of deployment of services by some third party software like VMware, Oracle etc. Thus it provides the fastest service to the users.

Reducing Infrastructure Cost

Virtualization essentially allows one computer to do the job of the multiple computers by means of sharing the resources of a single computer across multiple environments.

Virtual servers and virtual desktops allow hosting multiple operating systems and multiple applications locally and in remote locations. It lowers the expense by efficient use of the hardware resources.

It increases utilization rate for server and cost savings efficiently by altering the physical resources by virtual sharing.

Some other reasons are

1. To run old Apps
2. To access virus infected Data
3. To safely browse
4. Test software, upgrades or new configurations
5. To run Linux on top of Windows
6. To backup a entire operating system
7. To create a personal cloud computer
8. To reuse old hardware.

(OR)

19. (b) Write short notes on (i) Software virtualization (ii) network virtualization.

(i) Software virtualization

It is the virtualization of applications or computer programs. One of the most widely used software virtualization is Software Virtualization Solution (SVS) which is developed by Altris.

It is similar to hardware which is simulated as virtual machines. Software virtualization involves creating a virtual layer or virtual hard drive space where applications can be installed. From this virtual space, the application can be run as they have been installed onto host OS.

Once user finished using application, they can switch it off. When a application is switched off, any changes that the application made to the host OS will be completely reversed. This means that registry entries and installation directories will have no trace of the application being installed, executed at all.

Benefits of software virtualization are,

- The ability to run applications without making permanent registry or library changes.
- The ability to run multiple versions of the same application.
- The ability to install applications that would otherwise conflict with each other.

- The ability to test new applications in an isolated environment.
- It is easy to implement.

(ii) Network virtualization

Network virtualization is the process of combining hardware and software network resources and network functionality into a single, software based administrative entity which is said to be virtual network. Network virtualization involves platform virtualization. Network virtualization is categories into external network virtualization and internal network virtualization.

External network virtualization is combining of many networks into a virtual unit. Internal network virtualization is providing network like functionality to the software containers on a single system. Network virtualization enables connections between applications, services, dependencies and end users to be accurately emulated in the test environment.

20. (a) discuss the design considerations for storage network.

The best storage area network design for a customer will take into consideration a number of critical issues:

- Uptime and availability
- Capacity and scalability
- Security
- Replication and disaster recovery

Find out how each of these factors will influence storage area network design choices.

Uptime and availability

Because several servers will rely on a SAN for all of their data, it's important to make the system very reliable and eliminate any single points of failure. Most SAN hardware vendors offer redundancy within each unit like dual power supplies, internal controllers and emergency batteries.

In a typical storage area network design, each storage device connects to a switch that then connects to the servers that need to access the data. To make sure this path isn't a point of failure, the client should buy two switches for the SAN network. Each storage unit should connect to both switches, as should each server. If either path fails, software can failover to the other. Some programs will handle that failover automatically, but cheaper software may require you to enable the failover manually.

Capacity and scalability

A good storage area network design should not only accommodate the client's current storage needs, but it should also be scalable so that the client can upgrade the SAN as needed throughout the expected lifespan of the system.

Because a SAN's switch connects storage devices on one side and servers on the other, its number of ports can affect both storage capacity and speed. By allowing enough ports to support multiple, simultaneous connections to each server, switches can multiply the bandwidth to servers. On the storage device side, enough ports for redundant connections to existing storage units, as well as units to add later should be present.

Security

With several servers able to share the same physical hardware, the security plays an important role in a storage area network design.

Most of this security work is done at the SAN's switch level. Zoning allows giving only specific Servers access to certain LUNs, like a firewall allows communication on specific ports for a given IP address. If any outward-facing application needs to access the SAN, like a website, the switch should be configured so that only the server's IP address can access it.

Replication and disaster recovery

With so much data stored on a SAN, the client wants to build disaster recovery into the system. SANs can be set up to automatically mirror data to another site, which could be a failsafe SAN a few meters away or a disaster recovery (DR) site hundreds or thousands of miles away.

If client wants to build mirroring into the storage area network design, one of the first considerations is whether to replicate synchronously or asynchronously. Synchronous mirroring means that as data is written to the primary SAN, each

change is sent to the secondary and must be acknowledged before the next write can happen. The alternative is to asynchronously mirror changes to the secondary site. This replication can be configured to happen as quickly as every second, or every few minutes or hours. This means that the client could permanently lose some data, if the primary SAN goes down before it has a chance to copy its data to the secondary.

(OR)

20. (b) Explain in detail about object storage.

File systems or object storage

Object storage (also known as object-based storage) is a computer data storage architecture that manages data as objects, as opposed to other storage architectures like file systems which manage data as a file hierarchy and block storage which manages data as blocks within sectors and tracks. Each object typically includes the data itself, a variable amount of metadata, and a globally unique identifier.

Object storage can be implemented at multiple levels, including the device level (object storage device), the system level, and the interface level. In each case, object storage seeks to enable capabilities not addressed by other storage architectures, like interfaces that can be directly programmable by the application, a namespace that can span multiple instances of physical hardware, and data management functions like data replication and data distribution at object-level granularity. Object storage systems allow retention of massive amounts of unstructured data.

Object storage is used for purposes such as storing photos on Facebook, songs on Spotify, or files in online collaboration services, such as Dropbox. The majority of cloud storage available in the market uses the object storage architecture. Two notable examples are Amazon Web Services S3, which debuted in 2005, and Rackspace Files. Other major cloud storage services include IBM Bluemix, Microsoft Azure, Google Cloud Storage, Alibaba Cloud OSS, Oracle Elastic Storage Service and DreamHost based on Ceph.

Characteristics of Object Storage

- Performs best for big content and high storage throughput
- Data can be stored across multiple regions
- Scales infinitely to Petabytes (bigger than terabyte) and beyond
- Customizable metadata, not limited to number of tags

Advantages

- Scalable capacity
- Scalable performance
- Durable
- Low cost
- Simplified management
- Single Access Point
- No volumes to manage/resize/etc.

Disadvantages

- No random access to files
- The Application Programming Interface (API), along with command line shells and utility interfaces (POSIX utilities) do not work directly with object-storage
- Integration may require modification of application and workflow logic
- Typically, lower performance on a per-object basis than block storage

The Object Storage is suited for the following:

- Unstructured data
- Media (images, music, video)
- Web Content
- Documents
- Backups/Archives

- Archival and storage of structured and semi-structured data
- Databases
- Sensor data
- Log files

The Object Storage is not suited for the following:

- Relational Databases
- Data requiring random access/updates within objects

21. (a) write short notes on (i) Brokered cloud storage access. (ii) Storage location and tenancy.

(i) Brokered cloud storage access

Cloud Broker is an entity that manages the use, performance and delivery of cloud services, and relationships between cloud providers and cloud consumers.

All the data stored in the cloud. It can be located in the cloud service provider's system used to transfer data from sent and received. The cloud computing has no physical system that serves this purpose. To protect the cloud storage is the way to isolate data from client direct access. They are two services are created. One service for a broker with full access to storage but no access to the client, and another service for a proxy with no access to storage but access to both the client and broker. These important two services are in the direct data path between the client and data stored in the cloud. Under this system, when a client makes a request for data, here's what happens:

1. The request goes to the external service interface of the proxy.
2. The proxy using internal interface, forwards the request to the broker.
3. The broker requests the data from the cloud storage system.
4. The storage system returns the results to the broker.
5. The broker returns the results to the proxy.

The proxy completes the response by sending the data requested to the client.

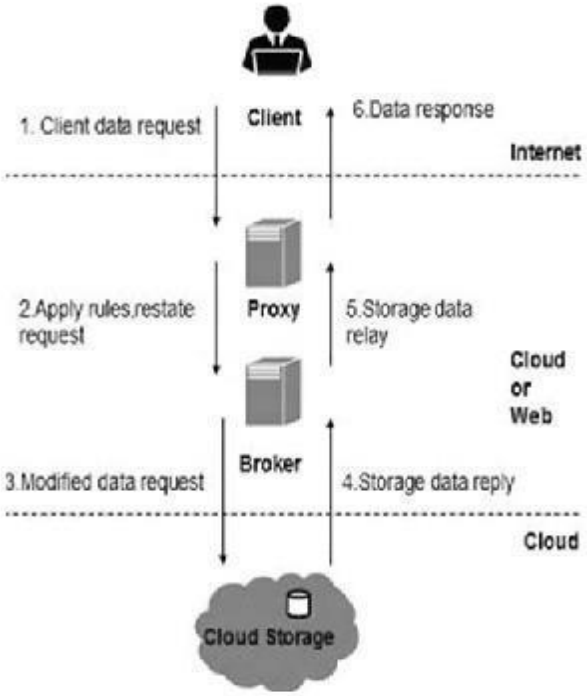

Fig 5.3 Cloud storage with proxy/broker service

Even if the proxy service is compromised, that service does not have access to the trusted key that is necessary to access the cloud storage. In the **multi-key solution**, not eliminated all internal service endpoints, but proxy service run at

a reduced trust level is eliminated. The creation of storage zones with associated encryption keys can further protect cloud storage from unauthorized access.

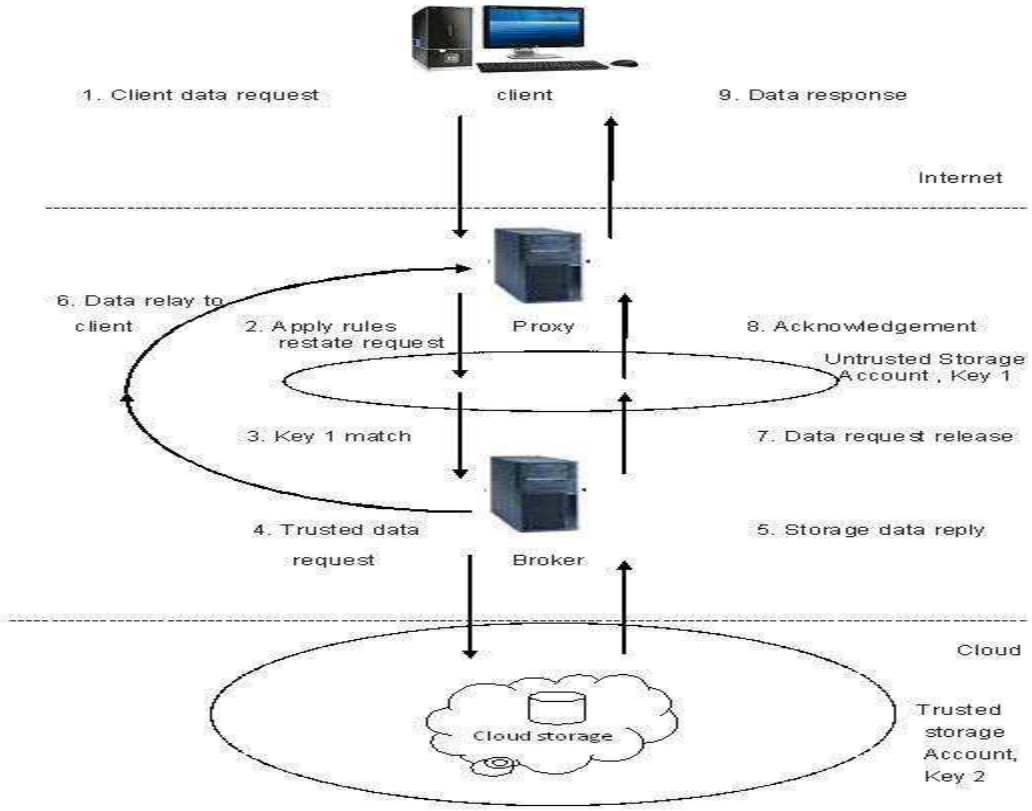

Storage zone with encrypted keys

Cloud brokers provide services in three categories:

Aggregation: A cloud broker combines and integrates multiple services into one or more new services.

Arbitrage: This is similar to service aggregation, except that the services being aggregated are not fixed.

Intermediation: The cloud broker give service by improving capability and providing value added services to cloud consumers. The improvement can be managing access to cloud services, identity management, performance reporting, enhanced security, etc.

Benefits of using a cloud broker

Benefits of using a cloud broker for a business or technical purpose include the following:

- Cloud interoperability - Integration between several cloud offerings.
- Cloud portability - Move application between different cloud vendors.
- Increase business continuity by reducing dependency from one cloud provider.
- Cost savings.

(ii) Storage location and tenancy

Cloud service providers as per their Service Level Agreements, need to contractually store and process data in locations that are predetermined by their contract. It gets the commitment for specific data site storage the cloud vendor is under contract to conform to privacy laws.

Because data stored in the cloud is usually stored from multiple tenants the each vendor has its own unique method for segregating one customer's data from another. It's important to understand how the specific service provider maintains data segregation. Cloud storage provider provides privileged access to storage. Most cloud service providers store data in an encrypted form to protect the data used in security mechanism. Hence, data cannot be accessed by the unauthorized user.

It is important to know what impact a disaster or interruption occur on the stored data. Since data are stored across multiples sites, it may not be possible to recover data in a timely manner.

(OR)

21. (b) (i) Explain virtualization security management. (ii)Explain briefly about virtual threats.

(i) Explain virtualization security management.

Historically, the development and implementation of new technology has preceded the full understanding of its inherent security risks, and virtualized systems are no different. The global adoption of virtualization is a relatively recent event, threats to the virtualized infrastructure.

A virtual machine (VM) is an operating system (OS) or application environment that is installed on software, which imitates dedicated hardware. The Virtual Machine (VM), Virtual Memory Manager (VMM), and hypervisor or host OS are the minimum set of components needed in a virtual environment.

Virtualization Types:

Based on the minimum set of components, we classify the Virtual Environments in the following distinct ways.

- Type 1 virtual environments are considered "full virtualization" environments and have VMs running on a hypervisor that interacts with the hardware.
- Type 2 virtual environments are also considered "full virtualization" but work with a host OS instead of a hypervisor.

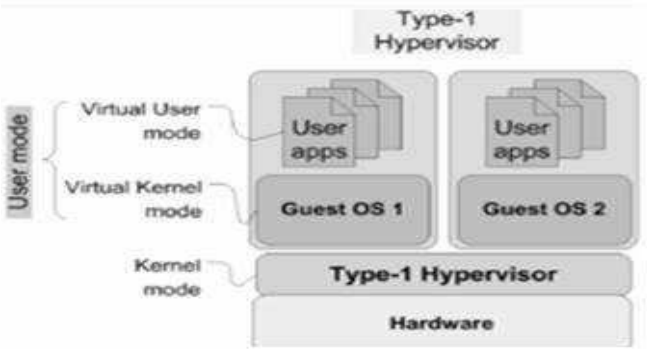

Figure 5.6a: Type 1

- Para virtualized environments offer performance gains by eliminating some of the emulation that occurs in full virtualization environments.
- Other type designations include hybrid virtual machines (HVMs) and hardware assisted techniques.

These classifications are somewhat ambiguous in the IT community at large. The most important thing to remember from a security perspective is that there is a more significant impact when a host OS with user applications and interfaces is running outside of a VM at a level lower than the other VMs (i.e., a Type 2 architecture). Because of its architecture, the Type 2 environment increases the potential risk of attacks against the host OS. For example, a laptop running VMware with a Linux VM on a Windows XP system inherits the attack surface of both OSs, plus the virtualization code (VMM).

Virtualization Management Roles:

The roles assumed by administrators are the Virtualization Server Administrator, Virtual Machine Administrator, and Guest Administrator. The roles assumed by administrators are configured in VMS and are defined to provide role responsibilities.

1. Virtual Server Administrator — This role is resp onsible for installing and configuring the ESX Server hardware, storage, physical and virtual networks, service console, and management applications.

2. Virtual Machine Administrator — This role is res ponsible for creating and configuring virtual machines, virtual networks, virtual machine resources, and security policies. The Virtual Machine Administrator creates, maintains, and provisions virtual machines.

3. Guest Administrator — This role is responsib le for managing a guest virtual machine Tasks typically performed by Guest Administrators include connecting virtual devices, adding system updates, and managing applications that may reside on the operating system.

(ii)Explain briefly about virtual threats.

Some threats to virtualized systems are general in nature, as they are inherent threats to all computerized systems (such as denial-of-service, or DoS, attacks). Other threats and vulnerabilities, however, are unique to virtual machines. Many VM vulnerabilities stem from the fact that vulnerability in one VM system can be exploited to attack other VM systems or the host systems, as multiple virtual machines share the same physical hardware.

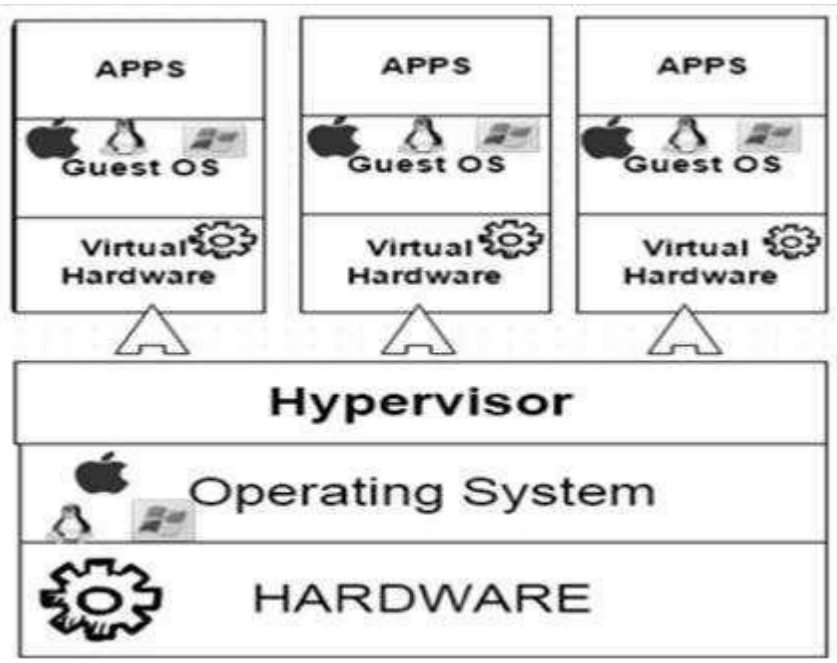

Figure 5.7: Virtual Threats

Some of the vulnerabilities exposed to any malicious-minded individuals regarding security in virtual environments:

Shared clipboard — Shared clipboard technology allows data to be tra nsferred between VMs and the host, providing a means of moving data between malicious programs in VMs of different security realms.

Keystroke logging — Some VM technologies enable the logging of keystr okes and screen updates to be passed across virtual terminals in the virtual machine, writing to host files and permitting the monitoring of encrypted terminal connections inside the VM.

VM monitoring from the host — because all network packets coming from or going to a VM pass through the host, the host may be able to affect the VM by the following:

1. Starting, stopping, pausing, and restart VMs.
2. Monitoring and configuring resources available to the VMs, including CPU, memory, disk, and network usage of VMs.
3. Adjusting the number of CPUs, amount of memory, amount and number of virtual disks and number of virtual network interfaces available to a VM.
4. Monitoring the applications running inside the VM.
5. Viewing, copying, and modifying data stored on the VM's virtual disks.

Virtual machine monitoring from another VM — Usually, VMs should not be able to directly access one another's virtual disks on the host.

Virtual machine backdoors — a backdoor, covert communications channel between the guest and host could allow intruders to perform potentially dangerous operations.

Prepared by,
CHITRA M [41205211**]**
Part Time / Guest Lecturer
120, Government Polytechnic College,
Purasaiwakkam, Chennai – 12.