

CLOUD COMPUTING

VI C.S.

INTRODUCTION

1

PREVIOUS YEARS QUESTIONS

PART-A

Q.1 What is cloud computing?

Ans. Cloud computing is a new age computer technology that is internet based. It is the next generation technology that utilizes the web based clouds to provide the services whenever users needs it.

Q.2 What are the essential things that must be followed before going to cloud computing platform?

Ans. Essential things that must be followed :

- Compliance issues
- Data storage types
- Maintaining data integrity in the cloud.
- Ensuring availability and access.
- Protection from loss of data
- Business continuity

Q.3 What do you mean by cloud in cloud computing?

Ans. Cloud is a combination of hardware, networks, services, storage and interface that helps in delivering computing as a service. It has three users :

1. End users
2. Business management users
3. Cloud service provider

Q.4 What are the different data types used in cloud computing?

Ans. There are different data types in cloud computing like emails, contracts, images, blogs etc. As we know that data is increasing day by day so it is example, if you want to store video then you need a new data type.

PART-B

Q.5 What are the risk in the migration into cloud? Also explain the process steps used in the migration into cloud.

[R.T.U. 2019]

What are the broad approaches for migration into cloud? Discuss the challenges and risks involved in this process.

[R.T.U. 2015]

Ans. There are three common approaches for migrating applications to the cloud:

- (i) Lift and shift
 - (ii) Refactoring
 - (iii) Extension or redesign
- Refactoring is the most preferred approach among the three since it involves making modest application code changes to ensure a smoother migration. The best approaches, though, are extension or heavily modifying code to fit the new cloud environment or a complete redesign to optimize the application for cloud.
- (i) **Lift and Shift** : Lift and shift means moving an application directly to a cloud host. In many cases, the

CLC2

approach works fine, but it won't take advantage of the cloud platform's full capabilities. But since re-architecting apps can be costly and time consuming, so some organizations prefer the lift and shift approach, which allows enterprises to take an inhouse app and replicate it in the cloud without modifying its design. However, because applications that are lifted and shifted to the cloud can't take full advantage of native cloud features, it's not always the most cost efficient migration approach.

That said, the lift and shift model still has its time and place. It's a good option, for example, for organizations that are "bleeding costs" from maintaining their own physical infrastructures.

For business critical applications that are poorly designed, there is significant risk in lifting and shifting them to the cloud. Without refactoring these applications will consume cloud resources inefficiently, thus generating a much higher public cloud bill and may even create performance and stability problems. In this case, given the importance of the application, it is well worth the investment to complete a partial or complete refactor in order to take full advantage of the cloud platform.

(ii) **Refactoring** : Refactoring is "the process of changing a software system in such a way that it does not alter the external behaviour of the code yet improves its internal structure". While refactoring can be applied to any programming language, the majority of refactoring current tools have been developed for the Java language.

One approach for refactoring is to improve the structure of source code at one point and then extend the same changes systematically to all applicable references throughout the program. The result is to make the code more efficient, scalable, maintainable or reusable, without actually changing any functions of the program itself.

Challenges and risks include the possibilities of introducing new bugs into the application and performance reduction, etc. Apart from that, it only takes advantage of some features of the cloud, as the code is not completely redesigned. Because of this, refactoring may incur a higher cloud bill that actually needed, whether it will be higher than that incurred by 'lift and shift' depends upon the application in question.

(iii) **Redesign/Extension** : This approach is about building applications from scratch in order to utilize the full potential of resources on cloud. When implemented with care, this approach brings the best utilization of cloud resources, but is very difficult to implement in itself.

Risks and challenges include much higher cost, since all the code must be rewritten in order to utilize the complete potential of the cloud resources. While doing that, there are high chances of introducing new bugs into

the code. Other than the above, the time to deployment is much slower than the other two.

Challenges and Risks Involved in Process : Cloud computing challenges have always been there. Companies are increasingly aware of the business value that cloud computing brings and are taking steps towards transition to the cloud. A smooth transition entails thorough understanding of the benefits as well as challenges involved. Like any new technology, the adoption of cloud computing is not free from issues. Some of the most important challenges are as follows :

1. **Security and Privacy** : The main challenge to cloud computing is how it addresses the security and privacy concerns of businesses thinking of adopting it. The fact that the valuable enterprise data will reside outside the corporate firewall raises serious concerns. Hacking and various attacks to cloud infrastructure would affect multiple clients even if only one site is attacked. These risks can be mitigated by using security applications, encrypted file systems, data loss software and buying security hardware to track unusual behavior across servers.

2. **Service Delivery and Billing** : It is difficult to assess the costs involved due to the on-demand nature of the services. Budgeting and assessment of the cost will be very difficult unless the provider has some good and comparable benchmarks to offer. The service-level agreements (SLAs) of the provider are not adequate to guarantee the availability and scalability. Businesses will be reluctant to switch to cloud without a strong service quality guarantee.

3. **Interoperability and Portability** : Businesses should have the leverage of migrating in and out of the cloud and switching providers whenever they want and there should be no lock-in period. Cloud computing services should have the capability to integrate smoothly with the on-premise IT.

4. **Reliability and Availability** : Cloud providers still lack round-the-clock service; this results in frequent outages. It is important to monitor the service being provided using internal or third-party tools. It is vital to have plans to supervise usage, SLAs, performance, robustness and business dependency of these services.

5. **Performance and Bandwidth Cost** : Businesses can save money on hardware but they have to spend more for the bandwidth. This can be a low cost for smaller applications but can be significantly high for the data-intensive applications. Delivering intensified complex data over the network requires sufficient bandwidth.

Because of this, many businesses are waiting for a reduced cost before switching to the cloud.

All these challenges should not be considered as road blocks in the pursuit of cloud computing. It is rather important to give serious consideration to these issues and the possible ways out before adopting the technology.

Q.6 What are the enabling technologies for cloud computing? Explain networking support for cloud computing. [R.T.U. 2019]

Ans. Enabling Technologies behind Cloud Computing : Cloud computing has been evolved by the advancement in various technologies e.g. distributed computing (cluster, grid computing etc.), internet technologies (Service-oriented architecture (SOA), web 3.0 etc.), hardware technologies (multi-core chips, virtualizations etc.) and system management technologies e.g. autonomic computing. Fig. shows the predecessor and the contributor technologies in the advent of cloud computing.




Fig. : The Contributor and Predecessor Technologies

1. SOA, Web 2.0, Web services, Mashups : Cloud computing is a paradigm shift from Application Oriented Architecture to Service Oriented Architecture (SOA). A Service Oriented Architecture is a set of collection of web services. Web services are self-contained, self-describing, platform independent programs that can be invoked over the internet. The orchestration of web services in a particular manner can be exposed as a single web service

to solve a particular activity over the internet. The advent of web 2.0 contributed in gluing up the web services together with the help of a new programming technique called AJAX (Asynchronous JavaScript and XML), REST (Representational State Transfer), RSS (Really Simple Syndication). AJAX is a technique to update the web pages without reloading the complete page. RSS distributes the up-to-date information to the web pages dynamically. The data exposed by the REST and RSS are extracted by a technique called a mashup. A mashup is defined as a website or web application that uses content or services from more than one source to create a completely new service.

2 Autonomic Computing : Autonomic Computing is a type of computing model in which the system is self-managing and adapt to the unpredictable changes. It is a collection of various existing technologies, including adaptive algorithms, human-computer interaction, machine learning algorithms, artificial intelligence algorithms, software agents, robotics, etc. autonomic computing means computing without or least human intervention. It is helpful to minimize the time required by the computer professional to resolve the system difficulties and the other maintenance work like software updates.

3 Utility and Grid Computing : The current model of cloud computing is an evolution of grid computing. Both the computing paradigms are based on utility computing. Utility computing is a service delivering model in which a service provider makes available the required resources to the customer and charges them for specific usage rather than a fixed rate. Cloud differs from grid computing in the sense that the cloud provides user-centric interfaces and does not require learning new commands and API as required in the case of grid computing.

4 Hardware Virtualization and Multi-core chips : Virtualization is a technique to create a virtual version of operating system, network, CPU, server, storage devices, etc. virtualization is an integrated solution to increase the resource utilization in a data center. Virtualization and cloud computing are used interchangeably but there is a significant difference between both the technologies, virtualization is a technology that manipulates hardware and cloud computing is a service that is a result of the manipulation done by virtualization. Virtualization is a basis for cloud computing that helps in improving resource utilization rate.

CLC.4

Networking Support : Cloud computing is a technique of resource sharing where servers and storage in multiple locations are connected by networks to create a pool of resources. Without networking to act as a backbone, the cloud is doomed. The support that networking provides to clouds is referred to as "networking support". When applications are run, resources are allocated from this pool of resources that are connected by networking and connected to the user as needed. The missions of connecting the resources (servers and storage) into a resource pool and then connecting users to the correct resources create the network's mission in cloud computing. For many cloud computing applications, network performance will be the key to cloud computing performance, this is how important networking support can get for a cloud.

Q.7 Explain Ethical issues in cloud computing.

[R.T.U. 2018/]

OR

What are the Ethical Issues in cloud computing?

[R.T.U. 2019, 2017/]

Ans. Ethical Issues : While the cloud may be flexible and cost-efficient, a lack of data safeguards and compliance standards makes security the largest hurdle to leap. As evidenced by the growing number of digital storage solutions that are being advertised in bar announcements and legal publications, not to mention to broader non-legal audiences, the cloud is here to stay. It presents a cost-effective storage solution, but also creates risks to the security of confidential client information. With proper due-diligence and reasonable care, attorneys can avoid the ethical pitfalls of cloud computing and experience all the benefits cloud computing has to offer for their firms and clients.

Security comprises of the followings :

- (i) Are the company's employees adequately screened, trained, etc.?
- (ii) Can the company's employees access documents?
- (iii) How secure is the electronic encryption, both in transmission and in storage?
- (iv) Are the physical premises hosting the data secure?
- (v) What happens if data is improperly accessed and what is the notification process?
- (vi) Who owns the premises, the servers and even the data?
- (vii) Where will data storage be located?
- (viii) How is data destruction handled?

The greatest threat in a cloud computing environment, according to CSA, is data loss, the prospect

B.Tech. (VI Sem.) CS Solved Papers

of seeing your valuable data disappear into the ether without a trace. A malicious hacker might delete a target's data out of spite but then, you could lose your data to a careless cloud service provider or a disaster, such as a fire, flood or earthquake. Compounding the challenge, encrypting your data to ward off theft can backfire if you lose your encryption key.

Other greatest cloud computing security risk is account or service traffic hijacking. Cloud computing adds a new threat to this landscape, according to CSA. If an attacker gains access to your credentials, he or she can eavesdrop on your activities and transactions, manipulate data, return falsified information and redirect your clients to illegitimate sites. "Your account or services instances may become a new base for the attacker. From here, they may leverage the power of your reputation to launch subsequent attacks," according to the report. As an example, CSA pointed to an XSS attack on Amazon in 2010 that let attackers hijack credentials to the site.

The key to defending against this threat is to protect credentials from being stolen. "Organizations should look to prohibit the sharing of account credentials between users and services and they should leverage strong two-factor authentication techniques where possible," according to CSA.

Q.8 Explain the types of cloud service development in detail.

[R.T.U. 2018/]

Ans. Services offered by Cloud Computing

1. **Storage as a service** (also known as disk space on demand), as you may expect, is the ability to leverage storage that physically exists at a remote site but is logically a local storage resource to any application that requires storage. This is the most primitive component of cloud computing and is a component or pattern that is leveraged by most of the other cloud computing components.
2. **Database as a service (DaaS)** provides the ability to leverage the services of a remotely hosted database; sharing it with other users and having it logically function as if the database were local. Different models are offered by different providers, but the power is to leverage database technology that would typically cost thousands of dollars in hardware and software licenses.
3. **Information as a service** is the ability to consume any type of information, remotely hosted, through a well defined interface such as an API. Examples include stock price information, address validation

local and context dependent. So, what is a cloud service provider to do? The answer is that a responsible approach to cloud computing will require an active attempt to engage with users and stakeholders to identify problems early. This requires an open mind and a willingness to respond to concerns. It also means recognising that cloud computing is not just a value-neutral tool. After all, it is in everybody's interest to work together to ensure that difficulties arising from the cloud's use do not outweigh the many potential advantages.

Q.12. *What is cloud computing? Give and explain the challenges, risk and approaches of migration into cloud.* [R.T.U. 2016]

Ans. Cloud Computing : According to Mell et al. (2011). "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models and four deployment models."

Main characteristics of cloud computing are :

On-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

Diverse dimensions can be adopted to classify cloud computing, two commonly used categories are : cloud deployment models and service models.

Challenges and Risks Involved in Process : Refer to Q.5.

CLOUD COMPUTING ARCHITECTURE 2

PREVIOUS YEARS QUESTIONS

PART-A

Q.1 Mention platforms which are used for large scale computing.

Ans. **Apache Hadoop** : It is an open source software platform for dispersed storage and distributed of huge data sets on compute bundles built from the product hardware. **MapReduce** : It enables the processing of massive datasets using cloud sources and other commodity hardware.

Q.2 What are the different layers that define cloud architecture?

Ans. The different layers used by cloud architecture :

1. Cluster or Cloud controller
2. Cluster controller
3. Walrus
4. NC or Node controller
5. SC or Storage controller

Q.3 Write the name of different models for deployment in cloud computing.

Ans. The different deployment models in cloud computing are :

1. Private cloud
2. Public cloud
3. Community cloud
4. Hybrid model.

Q.4 What are hybrid clouds?

Ans. These are combination of public clouds as well as private clouds. It is usually preferred over both the clouds because it applies the most healthy approach to the implementation of the cloud architecture. It includes the functionalities as well as features of both the worlds at the same time.

Q.5 What is private clouds?

Ans. Private clouds are used to keep the strategic operations and other reasons secure. It is a complete platform which is fully functional and can be owned, operated and restricted to only an organization or an industry. Now a days, most of the organizations have moved to private clouds due to security reasons.

PART-B

Q.6 What is need of data centers? Explain by providing case study of any industry where data center is used. [R.T.U. 2019, 2017]

Ans. **Need of data Centers** : A data center is actually a centralized repository which may be physical or virtual that houses computer systems for management, storage, and dissemination of information and data organized around a business entity. A data center can be classified generally into either Internet which supports few applications or Enterprise data center which is a custom shelf of applications.

Cloud Computing

The data center comprises:

1. **Design / Infrastructure** : Facility built to specific specifications and standards to meet the needs of today's high-tech hardware. With proper, cooling (HVAC), Generators, security systems and many other factors that can cost millions of dollars.
2. **Risk Management** : Your data is safe as it might reside in a multiple data center and regular backups are taken in case of disasters or crashes.
3. **Redundancy** : Most data centers have redundancies like power, cooling, bandwidth, and networking.
4. **Bandwidth** : A 100 Mbps at the office seems good but consider getting whooping 10 Gbps, only possible with a data center.
5. **Security** : Build to secure your data against threats or hacks. 24/7/365 security camera surveillance and staff on-site protecting the equipment.
6. **Compliance / Certifications** : Hardware is compliant and certified professionals manage all the data and the facility.
7. **Cost** : Creating and managing even a very small in-house data center can be a tedious task. Moreover, when you factor in costs then it requires massive amount which is not your responsibility when you take infrastructure solution.
8. **Scalability** : You can collocate your present servers or hosted hardware at a data center thus you can increase your space, power, and bandwidth within 24 hour
9. **Reduced Maintenance** : You don't need to manage anything or in infrastructure with weekly or monthly maintenance tasks. Save large sums of money

Case Study : Extending life and reducing costs of data center storage assets with Park Place Technologies
Dustin Jordan, Assistant Director of Technology Operations and Systems Management (TOSM) at the Texas Tech University System (TTUS), was faced with a dilemma. This relationship with Dell Inc., the OEM, was solid and the warranty maintenance for the servers and storage in his data centers were more than satisfactory. However, support for his legacy server, storage and data center infrastructure assets were more difficult to obtain. This data center is 350 miles from the nearest metropolitan area, and third-party service providers he had talked to could not supply both the cost-effectiveness and

CLC17

flexibility that he needed. To avoid replacing expensive devices prematurely, he needed extraordinary maintenance and support.

Texas Tech is a major comprehensive research university and medical institution that provides higher education to liberal arts, technical and medical students in several locations across Texas. Mr. Jordan's group of 14 full-time IT staff members supports hardware hosting and facilities management for servers and storage devices for the TTU system and its institutions.

TOSM's responsibilities include server and storage hosting, server management, local and remote server backups and database hosting and management. His group is also responsible for hosting the ERP system shared by Texas Tech University and the Texas Tech University Health Sciences Center, which includes the student, finance, financial aid and human resource systems for these institutions.

Q.7 Explain the working of MapReduce. [R.T.U. 2019]

OR

Write short note on MapReduce. [R.T.U. 2015]

Explain MapReduce model.

Ans. **MapReduce** : MapReduces programming model and associated implementation for processing large data sets. Programmer essentially just specifies two (sequential) functions: map and reduce. Program execution is automatically parallelized on large clusters of commodity PCs. MapReduce could be implemented on different architectures, but Google proposed it for clusters.

1. **A strategy or model for writing programs that can easily be made to process data in parallel.**
A framework that runs these programs in parallel, automatically handling the details of division of labor, distribution, synchronization and fault-tolerance. The model and the framework work together to make programs that are scalable, distributed and fault-tolerant. In the map-reduce programming model, work is divided into two phases: a map phase and a reduce phase. Both of these phases work on key-value pairs. What these pairs contain is completely up to us. They could be URLs paired with counts of how many pages link to them or movie IDs paired with ratings. It all depends on how we write and set up our map-reduce job.

A MapReduce program typically acts something like this:
1. Input data, such as a long text file, is split into key-value pairs. These key-value pairs are then fed to our mapper. (This is the job of the map-reduce framework.)

CLC.18

2. Our mapper processes each key-value pair individually and outputs one or more intermediate key-value pairs.
3. All intermediate key-value pairs are collected, sorted and grouped by key (again, the responsibility of the framework).
4. For each unique key, our reducer receives the key

B.Tech. (VI Sem.) CS Solved Papers

- with a list of all the values associated with it and aggregates these values in some way (adding them up, taking averages, finding the maximum, etc.) and outputs one or more output key-value pairs.
5. Output pairs are collected and stored in an output file (by the framework).

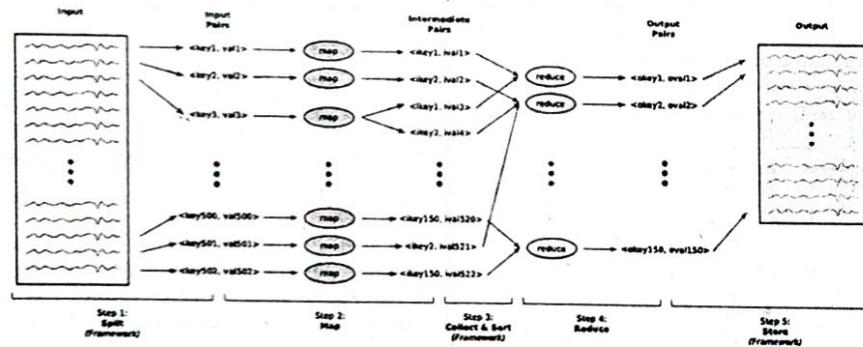


Fig. : Visualization of map-reduce

What makes this model so good for parallel programming should be apparent from the fig. Each key-value pair can be mapped or reduced independently. This means that many different processors or even machines, can each take a section of the data and process it separately, a classic example of data parallelism. The only real step where synchronization is needed is during the collecting and sorting phase, which can be handled by the framework (and, when done carefully, even this can be parallelized). So, when we can fit a problem into this model, it can make parallelization very easy. What may seem less obvious is how a problem can be solved with this model in the first place.

Programming Model


- (i) Transforms set of input key-value pairs to set of output values (notice small modification compared to paper) list (k2, v2) →
- (ii) Map: (k1, v1) →
- (iii) MapReduce library groups all intermediate pairs with same key together list (k3, v3) →
- (iv) Reduce: (k2, list (v2)) – Usually, zero or one output value per group – Intermediate values supplied via iterator (to handle lists that do not fit in memory).

Q.8 Explain the architecture of cloud computing in detail. [R.T.U. 2018]

Ans. Cloud Computing Architecture : Cloud computing architecture comprises of many cloud components, which are loosely coupled. Cloud architecture divided into two parts :

- Front End
- Back End

Each of the ends is connected through a network, usually Internet. The following diagram shows the graphical view of cloud computing architecture:

**Cloud Computing**

Front End : The front end refers to the client part of cloud computing system. It consists of interfaces and applications that are required to access the cloud computing platforms. Example : Web Browser.

Back End : The back end refers to the cloud itself. It consists of all the resources required to provide cloud computing services. It comprises of huge data storage, virtual machines, security mechanism, services, deployment models, servers, etc.

Q.9 Explain models of cloud computing in detail.

[R.T.U. 2018]

Ans. Cloud Computing Model

Single Server : Single server cloud templates embody the use of one server, physical or virtual, that includes a web server, an application, and a database. Single server architectures are not widely used, as they have potential security risks which can compromise the whole system. These architectures are usually deployed for development purposes, enabling developers to strengthen functionality without having to deal with connectivity and communication issues between different servers.

Single Site : Single-site platforms utilize the single server architecture and divide the layers in different systems thus modeling a three-tier architecture. As the system is administered in the same location, single-site architecture is maintained as up and running.

Redundant three-tier Model : Redundant three-tier architectures have the capability of appending another set of the same components to look after redundancy. These pairs of components increase the complexity of the system but are required to tackle failover and manage recovery protection. Constructing redundant infrastructures requires a well-executed plan for the components within each layer for horizontal scaling, along with a plan for how the traffic will circulate from one layer to another meant for vertical scaling.

Auto-scaling Model : The main takeaway of cloud computing is the ability to utilize a particular service or resource when it is needed. This autoscaling feature enables the cloud to scale horizontally which is to shrink or grow the number of running server instances in a largely coupled system with changes in demand of application over time.

Different sorts of architectures are supplied to the clients. It is the responsibility of the client to make sure that the model he prefers for his business must suit his business needs and requirements in the long run. At times, it is necessary to gain information regarding cloud from

aggregate data from thousands of small businesses. Each business that uses a cloud service increases the value of that service as a potential target. This concentrates risk on a single point of failure. A disaster at a cloud provider can affect its every customer.

2. **Security Risks at the Vendor :** When a cloud service vendor supplies a critical service for your business and stores critical data, such as customer payment data and your mailing lists.
3. **Compliance and Legal Risks :** Many data security regulations are intended to protect a specific type of data. For example, HIPAA requires healthcare providers to protect patient data. PCI DSS requires anyone who accepts credit cards to protect cardholder data. The companies which are covered by these regulations not only required to protect the data. But they are also typically required to know where the data resides, who is allowed to access it and how it is protected.
4. **Risks Related to Lack of Control :** When we host and maintain a service on a local network, then we have complete control over the features we choose to use. If we want to change the service in the future, we are in control. However, when we use a cloud service provider, the vendor is in control. We have no guarantee that the features we use today will be provided for the same price tomorrow. The vendor can double its price, and if our clients are depending on that service, then we might be forced to pay. Also, who controls access to our data in a cloud service? What happens if we are not able to make payment?
5. **Risks Related to Availability :** No service can guarantee 100% uptime. When you rely on a cloud service for a business-critical task, then you are putting the viability of your business in the hands of two services: the cloud vendor and your ISP. If your internet access goes down, then it will take your vendor's cloud service with it. If we need the cloud service to process customer payments or access important data, then you have to wait until the internet is back up.

Q.12 Describe the following service delivery models by giving the suitable examples in industry :

- (i) IaaS
- (ii) SaaS

[R.T.U. 2015]

Ans. (i) IaaS : Infrastructure as a service (IaaS) is actually data center as a service or the ability to remotely access computing resources. In essence, we lease a physical server that is yours to do with as we will and, for all practical purposes, is your data center or at least part of a data center. The difference with this approach versus more mainstream cloud computing is that instead of using an interface and a metered service, we have access to the entire machine and the software on that machine. In short, it is less packaged.

Example : Amazon Web Services, Cisco Metapod, Microsoft Azure, etc.

(ii) SaaS : Cloud application services or Software as a Service (SaaS), represent the largest cloud market and are still growing quickly. SaaS uses the web to deliver applications that are managed by a third party vendor and whose interface is accessed on the clients' side. Most SaaS applications can be run directly from a web browser without any downloads or installations required, although some require plugins.

Because of the web delivery model, SaaS eliminates the need to install and run applications on individual computers. With SaaS, it's easy for enterprises to streamline their maintenance and support, because everything can be managed by vendors: applications, runtime, data, middleware, OSes, virtualization, servers, storage and networking.

Popular SaaS offering types include email and collaboration, customer relationship management and healthcare related applications. Some large enterprises that are not traditionally thought of as software vendors have started building SaaS as an additional source of revenue in order to gain a competitive advantage.

Example : Google Apps, Salesforce, Workday, Concur, Citrix GoToMeeting, Cisco WebEx.

Q.13 Explain cloud ecosystem.

Ans. Cloud Ecosystem: With the emergence of various Internet clouds, an ecosystem of providers, users and technologies has appeared. This ecosystem has evolved public clouds. Strong interest is growing in open source cloud computing tools that let organizations build their own IaaS clouds using their internal infrastructures. Private and hybrid clouds are not exclusive, since public clouds are involved in both cloud types. A private/hybrid cloud allows remote access to its resources over the internet using remote web service interfaces such as that used in Amazon EC2.

An ecosystem was suggested by Sotomayor, et al. (fig.) for building private clouds. They suggested four

Q.15 Define public private and hybrid cloud.

OR

Differentiate between public, private and hybrid cloud according to their functionality.

[R.T.U. 2019, 2017]

Ans. Cloud computing comes in three forms: public clouds, private clouds and hybrids clouds. Depending on the type of data you're working with, you'll want to compare public, private and hybrid clouds in terms of the different levels of security and management required.




Fig.

Public Clouds : A public cloud is one in which the services and infrastructure are provided off-site over the internet. These clouds offer the greatest level of efficiency in shared resources; however, they are also more vulnerable than private clouds. A public cloud is the obvious choice when

- Your standardized workload for applications is used by lots of people, such as e-mail.
- You need to test and develop application code.
- You have SaaS (Software as a Service) applications from a vendor who has a well-implemented security strategy.
- You need incremental capacity (the ability to add computer capacity for peak times).
- You're doing collaboration projects.
- You're doing an ad-hoc software development project using a Platform as a Service (PaaS) offering cloud.

Many IT department executives are concerned about public cloud security and reliability. Take extra time to ensure that you have security and governance issues well planned or the short-term cost savings could turn into a long-term nightmare.

Private Clouds : A private cloud is one in which the services and infrastructure are maintained on a private network. These clouds offer the greatest level of security and control, but they require the company to still purchase

and maintain all the software and infrastructure, which reduces the cost savings. A private cloud is the obvious choice when

- Your business is your data and your applications. Therefore, control and security are paramount.
- Your business is part of an industry that must conform to strict security and data privacy issues.
- Your company is large enough to run a next generation cloud data center efficiently and effectively on its own.

To complicate things, the lines between private and public clouds are blurring. For example, some public cloud companies are now offering private versions of their public clouds. Some companies that only offered private cloud technologies are now offering public versions of those same capabilities.

Hybrid Clouds : A hybrid cloud includes a variety of public and private options with multiple providers. By spreading things out over a hybrid cloud, you keep each aspect at your business in the most efficient environment possible. The downside is that you have to keep track of multiple different security platforms and ensure that all aspects of your business can communicate with each other. Here are a couple of situations where a hybrid environment is best.

- Your company wants to use a SaaS application but is concerned about security. Your SaaS vendor can create a private cloud just for your company inside their firewall. They provide you with a virtual private network (VPN) for additional security.
- Your company offers services that are tailored for different vertical markets. You can use a public cloud to interact with the clients but keep their data secured within a private cloud.

The management requirements of cloud computing become much more complex when you need to manage private, public and traditional data centers all together. You'll need to add capabilities for federating these environments.

Q.16 Explain various service layers in layered architecture of cloud with the help of a neat and labelled diagram.

[R.T.U. 2019]

OR

Explain various service layers in layered architecture of cloud with the help of suitable examples.

[R.T.U. 2015]

OR

Explain cloud deployment model and layers of cloud computing architecture.

CLC-24

Ans. Cloud Deployment Models : Most organizations focusing on leveraging the cloud in order to cut capital expenditure and control operating costs, there is aggressive growth in business for cloud adoption. However, the cloud can bring security risks and challenges for IT Management, which can be more expensive for the organization to deal with, even considering the cost saving achieved by moving to the cloud. Therefore, it is very important for businesses to understand their requirements before opting for various deployment models available on the cloud. There are primarily four cloud deployment models, which are discussed in further sections, along with scenarios in which a business could opt for each. These models have been recommended by the National Institute of Standards and Technology (NIST).

1. **The Private Cloud:** This model doesn't bring much in terms of cost efficiency; it is comparable to buying, building and managing our own infrastructure. Still, it brings in tremendous value from a security point of view. During their initial adaptation to the cloud, many organizations face challenges and have concerns related to data security. These concerns are taken care of by this model, in which hosting is built and maintained for a specific client. The infrastructure required for hosting can be on-premises or at a third-party location.

Security concerns are addressed through secure access VPN for the physical location within the client's firewall system.

Furthermore, for mission-critical applications we need to consider downtime in terms of internet availability, quality and performance. Hence, hosting the application with an on-premises private cloud is the suggested approach.

In addition to security reasons, this model is adopted by organizations in cases where data or applications are required to conform to various regulatory standards such as SOX, HIPAA or SAS 70, which may require data to be managed for privacy and audits that govern the corporation. For example, for the healthcare and pharmaceutical industries, moving data to the cloud may violate the norms. Similarly, different countries have different laws and regulations for managing and handling data, which can interrupt the business if cloud is under different jurisdiction.

Several SaaS Applications, such as Sugar CRM, provide options to their clients to maintain their data on their own premises to ensure data privacy is maintained according to the requirements of the particular business. Amazon also provides the option of a virtual private cloud.

2. **The Public Cloud :** The public cloud deployment model represents true cloud hosting. In this

B.Tech. IV Sem./CS Solved Papers

deployment model, services and infrastructure are provided to various clients. Google is an example of a public cloud. This service can be provided by a vendor free of charge or on the basis of a pay-per-user license policy.

This model is best suited for business requirements wherein it is required to manage load spikes, host SaaS applications, utilize interim infrastructure for developing and testing applications and manage applications which are consumed by many users that would otherwise require large investment in infrastructure from businesses.

This model helps to reduce capital expenditure and bring down operational IT costs.

3. **The Hybrid Cloud :** This deployment model helps businesses to take advantage of secured applications and data hosting on a private cloud, while still enjoying cost benefits by keeping shared data and applications on the public cloud. This model is also used for handling cloud bursting, which refers to a scenario where the existing private cloud infrastructure is not able to handle load spikes and requires a fallback option to support the load. Hence, the cloud migrates workloads between public and private hosting without any inconvenience to the users.

Many PaaS deployments expose their APIs, which can be further integrated with internal applications or applications hosted on a private cloud, while still maintaining the security aspects. Microsoft Azure and Force.com are two examples of this model.

4. **The Community Cloud :** In the community deployment model, the cloud infrastructure is shared by several organizations with the same policy and compliance considerations. This helps to further reduce cost as compared to a private cloud, as it is shared by larger group. Various state-level government departments requiring access to the same data relating to the local population or information related to infrastructure, such as hospitals, roads, electrical stations, etc., can utilize a community cloud to manage applications and data.




Fig. 1

Cloud Computing

Cloud computing is not a "silver-bullet" technology; hence, investment in any deployment model should be made based on business requirements, the criticality of the application and the level of support required.

Layered Cloud Architectural Development

The architecture of a cloud is developed at three layers: infrastructure, platform and application as demonstrated in Fig. 2. These three development layers are implemented with virtualization and standardization of hardware and software resources provisioned in the cloud. The services to public, private and hybrid clouds are conveyed to users through the networking support over the Internet and Intranet involved. It is clear that the infrastructure layer is deployed first to support IaaS type of services. This infrastructure layer serves as the foundation to build the platform layer of the cloud for supporting PaaS services. In turn, the platform layer is a foundation to implement the application layer for SaaS applications. Different types of cloud services demand to apply the resources, separately.

The infrastructure layer is built with virtualized compute, storage and network resources. The abstraction of these hardware resources is meant to provide the flexibility demanded by users. Internally, the virtualization realizes the automated provisioning of resources and optimizes the infrastructure management process. The platform layer is for general-purpose and repeated usage of the collection of software resources. This layer provides the users with an environment to develop their applications, to test the operation flows and to monitor the execution results and performance. The platform should be able to assure the users with scalability, dependability and security protection. In a way, the virtualized cloud platform serves as a "system middleware" between the infrastructure and application layers of the cloud. The application layer is formed with a collection of all needed software modules for SaaS applications. Service applications in this layer include daily office management work, such as information retrieval, document, processing, calendar and authentication services, etc.

The application layer is also heavily used by enterprises in business marketing and sales, consumer relationship management (CRM), financial transactions, supply chain management etc. It should be noted that not all cloud services are restricted to a single layer. Many applications may apply resources at mixed layers. After all, the three layers are built from bottom up with a dependence relationship. From the provider's perspective, the services at various layers demand different amounts of function, support and resource management by the providers. In general, the SaaS demands the most work

CLC-25

from the provider, the PaaS in the middle and IaaS the least. For an example, Amazon EC2 provides not only virtualized CPU resources to users but also the management of these provisioned resources. Services at the application layer demands more work from the providers. The best example is the Salesforce CRM service in which the provider supplies not only the hardware at the bottom layer and the software at the top layer, but also provides the platform and software tools for user application development and monitoring. Layered architectural development of the cloud platform for IaaS, PaaS and SaaS applications over the Internet and Intranet.

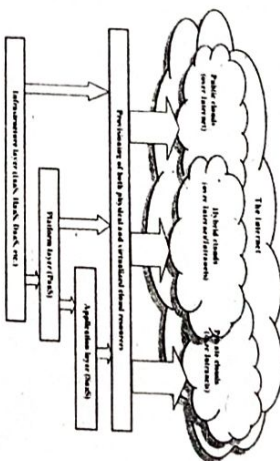


Fig. 2

Q.17 What are the high level languages that can be used in cloud computing programming? Explain.

[R.T.U. 2017]

Ans. A high-level language (HLL) is a programming language such as C, FORTRAN, or Pascal that enables a programmer to write programs that are more or less independent of a particular type of computer. Such languages are considered high-level because they are closer to human languages and further from machine languages.

There is no direct relationship between the characteristics of programming languages and the concepts involved in today's definition of cloud computing. Cloud computing platforms themselves are built in many languages from C through to Python, Java, C++, Ruby.

Each aspect of cloud computing has languages which are suitable to the task. C++ is better for implementing low level cloud management layers which allocate processor resources and do virtualization. Many conventional web languages from php through to Python, Ruby, and Scala are useful running under frameworks like Apache or WSGI. NodeJS (which is a framework paradigm, not a language), along with Twisted or Tornado

CLC.30

2. Public Cloud Storage : Public cloud storage is where the enterprise and storage service provider are separate and there aren't any cloud resources stored in the enterprise's data center. The cloud storage provider fully manages the enterprise's public cloud storage.

3. Private Cloud Storage : A form of cloud storage where the enterprise and cloud storage provider are integrated in the enterprise's data center. In private cloud storage, the storage provider has infrastructure in the enterprise's data center that is typically managed by the storage provider. Private cloud storage helps resolving the potential for security and performance concerns while still offering the advantages of cloud storage.

4. Hybrid Cloud Storage : Hybrid cloud storage is a combination of public and private cloud storage where some critical data resides in the enterprise's private cloud while other data is stored and accessible from a public cloud storage provider.

Q.20 Write short note on Hadoop.

[R.T.U. 2015]

OR

Explain hadoop and its architecture.

Ans. Hadoop : Hadoop is an open-source software framework for storage and large-scale processing of datasets on clusters of commodity hardware. There are mainly five building blocks inside this runtime environment (from bottom to top).




Fig. 1 : Hadoop architecture

1. The **cluster** is the set of host machines (**nodes**). Nodes may be partitioned in **racks**. This is the hardware part of the infrastructure.
2. The **YARN Infrastructure** (Yet Another Resource Negotiator) is the framework responsible for providing the computational resources (e.g., CPUs, memory, etc.) needed for application executions. Two important elements are:
3. The **Resource Manager** (one per cluster) is the

B.Tech. (VI Sem.) CS Solved Papers

master. It knows where the slaves are located (Rack Awareness) and how many resources they have. It runs several services, the most important is the **Resource Scheduler** which decides how to assign the resources.




Fig. 2

4. The **Node Manager** (many per cluster) is the slave of the infrastructure. When it starts, it announces himself to the Resource Manager. Periodically, it sends an heartbeat to the Resource Manager. Each Node Manager offers some resources to the cluster. Its resource capacity is the amount of memory and the number of vcores. At run-time, the Resource Scheduler will decide how to use this capacity: a **Container** is a fraction of the NM capacity and it is used by the client for running a program.




Fig. 3

5. The **HDFS Federation** is the framework responsible for providing permanent, reliable and distributed storage. This is typically used for storing inputs and output (but not intermediate ones).
6. Other alternative storage solutions. For instance, Amazon uses the Simple Storage Service (S3).
7. The **MapReduce Framework** is the software layer implementing the **MapReduce paradigm**. The YARN infrastructure and the HDFS federation are completely decoupled and independent: the first one provides resources for running an application while the second one provides storage. The MapReduce framework is only one of many possible framework which runs on top of YARN (although currently is the only one implemented).

Cloud Computing

YARN: Application Startup




Fig. 4

In YARN, there are at least three actors:

1. The Job Submitter (the client)
 2. The Resource Manager (the master)
 3. The Node Manager (the slave)
- The application startup process is the following:
1. A client submits an application to the Resource Manager.
 2. The Resource Manager allocates a container.
 3. The Resource Manager contacts the related Node Manager.
 4. The Node Manager launches the container.
 5. The Container executes the application master.

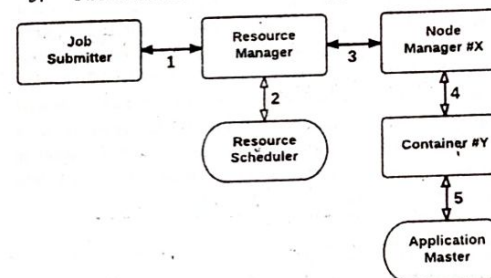


Fig. 5

The application master is responsible for the execution of a single application. It asks for containers to the Resource Scheduler (Resource Manager) and executes specific programs (e.g., the main of a Java class) on the obtained containers. The application master knows the application logic and thus it is framework-specific. The MapReduce framework provides its own implementation of an application master.

The resource manager is a single point of failure in YARN. Using application masters, YARN is spreading over the cluster the metadata related to running applications. This reduces the load of the resource manager and makes it fast recoverable.

Q.21 What do you understand by the term parallel and distributed computing environments? Also define parallel programming tools.

PART-A

Q.1 *Explain the differences between emulation, native virtualization and host virtualization?*


[R.T.U. 2019, 2017]

Ans. Difference between emulation, native and host virtualization : For many, emulation and virtualization go hand in hand, but there are actually some really key differences. When a device is being emulated, a software-based construct has replaced a hardware component. It's possible to run a complete virtual machine on an emulated server. However, virtualization makes it possible for that virtual machine to run directly on the underlying hardware, without needing to impose an emulation tax (the processing cycles needed to emulate the hardware).

Q.2 *What is the requirement of virtualization platform in implementing cloud?*

Ans. The requirement of virtualization platform in implementing cloud is to :

- (a) Manage the service level policies
- (b) Cloud operating system
- (c) Virtualization platforms helps to keep the backend level and user level concepts different from each other.

 **Q.3** *Define virtualization.*

Ans. Virtualization : Virtualization is software that separates physical infrastructures to create various

dedicated resources. It is the fundamental technology that powers cloud computing.

Q.4 *What are the types of storage virtualization?*

Ans. Types of storage virtualization are :

- Hardware assisted virtualization
- Kernel level virtualization
- Hypervisor virtualization
- Para-virtualization
- Full virtualization

Q.5 *Write the concept of file based storage virtualization.*

Ans. This type of virtualization is used for specific purpose and can apply to Network Attached Storage (NAS) system. File based storage virtualization in cloud computing utilizes server message block or network file system protocols and with its help of it breaks the dependency in a normal network attached storage array.

PART-B

 **Q.6** *Explain the virtualization of CPU in details.*

[R.T.U. 2019]

Ans. CPU Virtualization : 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. The critical instructions are divided into three categories: privileged instructions, control-sensitive instructions, and behavior-sensitive instructions. Privileged instructions execute in a privileged mode and will be trapped if executed outside this mode. Control-sensitive instructions attempt to change the configuration of resources used. Behavior-sensitive instructions have different behaviors depending on the configuration of resources, including the load and store operations over the virtual memory.

A CPU architecture is virtualizable if it supports the ability to run the VMs privileged and unprivileged instructions in the CPU's user mode while the VMM runs in supervisor mode. When the privileged instructions including control and behavior-sensitive instructions of a VM are executed, they are trapped in the VMM. In this case, the VMM acts as a unified mediator for hardware access from different VMs to guarantee the correctness and stability of the whole system. However, not all CPU architectures are virtualizable. RISC CPU architectures can be naturally virtualized because all control and behavior-sensitive instructions are privileged instructions. On the contrary, x86 CPU architectures are not primarily designed to support virtualization. This is because about 10 sensitive instructions, such as SGDT and SMSW, are not privileged instructions. When these instructions execute in virtualization, they cannot be trapped in the VMM.

On a native UNIX-like system, a system call triggers the 80h interrupt and passes control to the OS kernel. The interrupt handler in the kernel is then invoked to process the system call. On a para-virtualization system such as Xen, a system call in the guest OS first triggers the 80h interrupt normally. Almost at the same time, the 82h interrupt in the hypervisor is triggered. Incidentally, control is passed on to the hypervisor as well. When the hypervisor completes its task for the guest OS system call, it passes control back to the guest OS kernel. Certainly, the guest OS kernel may also invoke the hypercall while it's running. Although paravirtualization of a CPU lets unmodified applications run in the VM, it causes a small performance penalty.

Hardware Assisted CPU Virtualization : This technique attempts to simplify virtualization because full or paravirtualization is complicated. Intel and AMD add an additional mode called privilege mode level (some people

call it Ring-1) to x86 processors. Therefore, operating systems can still run at Ring 0 and the hypervisor can run at Ring -1. All the privileged and sensitive instructions are trapped in the hypervisor automatically. This technique removes the difficulty of implementing binary translation of full virtualization. It also lets the operating system run in VMs without modification.

Example

Intel Hardware Assisted CPU Virtualization : Although x86 processors are not virtualizable primarily, great effort is taken to virtualize them. They are used widely in comparing RISC processors that the bulk of x86-based legacy systems cannot discard easily. Virtualization of x86 processors is detailed in the following sections. Intel's VT-x technology is an example of hardware-assisted virtualization, as shown in Fig. Intel calls the privilege level of x86 processors the VMX Root Mode. In order to control the start and stop of a VM and allocate a memory page to maintain the CPU state for VMs, a set of additional instructions is added. At the time of this writing, Xen, VMware, and the Microsoft Virtual PC all implement their hypervisors by using the VT-x technology.




Fig. : Intel hardware assisted CPU virtualization

Generally, hardware assisted virtualization should have high efficiency. However, since the transition from the hypervisor to the guest OS incurs high overhead switches between processor modes, it sometimes cannot outperform binary translation. Hence, virtualization systems such as VMware now use a hybrid approach, in which a few tasks are offloaded to the hardware but the rest is still done in software.

Q.7 Write short note on KVM.

[R.T.U. 2018]

Ans. KVM : KVM represents a full hardware virtualisation platform with loadable kernel modules, giving users the freedom to run a range of Linux distros with any kernel. Each virtual machine has private virtualised hardware including network card, disk and graphics adapter, and with no possibility of overselling, you get guaranteed resources at your disposal any time day or night!

KVM In my opinion is the up and comer in the Virtualization world. It is back by Redhat and they have removed support for Xen by default on Enterprise Linux 6, in favor of KVM. It has a lot of features and is similar in certain ways to Xen, as in it supports Para virtualization via the Virtio IO framework. Our cloud servers are based on KVM virtualization

KVM is a type-2 hypervisor built into the Linux kernel as a module and will ship with any Linux distribution moving forward as no work is required for the Linux distributions to add KVM. Having a virtualization platform built-in to the Linux kernel will be valuable to many customers looking for virtualization within a Linux based infrastructure; however these customers will lose the flexibility to run a bare-metal hypervisor, configure the hypervisor independent of the host operating system, and provide machine level security as a guest can bring down the operating system on KVM.

Q.8 Write short note on memory.

[R.T.U. 2018]

Ans.(b) Memory : In the past, memory was limited by the particular device in question. Ran out of memory need a USB drive to backup current device. Cloud computing provides increased storage, so it won't have to be worry about running out of space on hard drive.

The software is already installed online, so it won't need to be install it. There are numerous cloud computing applications available for free, such as Dropbox, and increasing storage size and memory is affordable.

In-memory computing is the storage of information in the main random access memory (RAM) of dedicated servers rather than in complicated relational databases operating on comparatively slow disk drives. In-memory computing helps business customers, including retailers, banks and utilities, to quickly detect patterns, analyze massive data volumes on the fly, and perform their

centers on the same physical infrastructure, which can simultaneously be used by separate applications and organizations. This not only helps in optimal IT infrastructure and resource utilization, but also in reducing data center capital and operational costs. Virtualization of devices such as server, storage device, network and even an operating system has induced the concept a Software defined datacenter or SSDC

Hardware is mostly the costliest asset of the Data centers. The perks of virtualisation are even higher as it is easy to maintain, consumes less electricity and there are even lesser occasions of downtime. Reducing the cost of hardware is directly proportional to reduced cost. One can have total backup of virtual servers along with the backups and snapshots of virtual machines. The virtual machines can be locomoted from one server to another as well as redeployed fast and super-easily.

Q.11 What is network virtualization? Describe the various components in network virtualization.

[R.T.U. 2018, 2015]

Ans. Network virtualization is a method of combining the available resources in a network by splitting up the available bandwidth into channels, each of which is independent from the others and each of which can be assigned (or reassigned) to a particular server or device in real time. Each channel is independently secured. Every subscriber has shared access to all the resources on the network from a single computer.

Network management can be a tedious and time consuming business for a human administrator. Network virtualization is intended to improve productivity, efficiency and job satisfaction of the administrator by performing many of these tasks automatically, thereby disguising the true complexity of the network. Files, images, programs and folders can be centrally managed from a single physical site. Storage media such as hard drives and tape drives can be easily added or reassigned. Storage space can be shared or reallocated among the servers.

Network virtualization is intended to optimize network speed, reliability, flexibility, scalability and security. Network virtualization is said to be especially effective in networks that experience sudden, large and unforeseen surges in usage.

The following are the basic components of network virtualization:

1. Virtual network interface cards (VNICs):

VNICs are virtual network devices with the same datalink interfaces as a physical NIC. You configure VNICs over an underlying datalink. When VNICs are configured, they

behave like physical NICs. In addition, the system's resources treat VNICs as if they were physical NICs. A VNIC has an automatically generated MAC address. Depending on the network interface in use, you can explicitly assign to a VNIC a MAC address other than this default address.

2. Virtual switches: When you create a VNIC, a virtual switch is automatically created. In accordance with Ethernet design, if a switch port receives an outgoing packet from the host connected to that port, that packet cannot go to a destination on the same port. This design is a drawback for systems that are configured with virtual networks because the virtual networks share the same NIC. The outgoing packets go through a switch port out onto the external network. The incoming packets cannot reach their destination zone because the packets cannot return through the same port that they were sent through. Virtual switches provide these zones with a method to pass packets. The virtual switch opens a data path for the virtual networks to communicate with one another.

3. Etherstubs: Etherstubs are pseudo Ethernet NICs. You can create VNICs over etherstubs instead of over physical links. VNICs over an etherstub become independent of the physical NICs on the system. With etherstubs, you can construct a private virtual network that is isolated both from the other virtual networks on the system and from the external network. For example, if you want to create a network environment whose access is limited only to your company developers and not to the network at large, etherstubs can be used to create such an environment.

Q.12 What is the difference between KVM and Xen?

[R.T.U. 2017]

Ans. Difference between KVM and Xen:

Xen : Xen is a bare metal hypervisor, which makes it capable of running multiple instances of virtual machines on a single host. These hosts are not constrained to the kernel of the host and for that matter do not even have to run Linux in the VPS. Xen Virtualization is capable of hosting Windows and BSD operating systems as seamlessly as a Linux guest. Additionally Xen is a very light hypervisor with a small foot print. This leaves the valuable resources you need where you need it, for the guest virtual machines.

Another great benefit from a consumer perspective is that Xen cannot be over-subscribed. Each guest's resources are allocated all the time on the host node. No sharing memory and hoping some is there when you need it! Xen supports both Hardware Virtual Machine (HVM) and Paravirtualization (PV) in the hypervisor

Resource management : You have to effectively manage VMs running on a mass of physical computing nodes (called virtual clusters) in a high-performance virtualized computing environment. This involves virtual cluster deployment, monitoring and management over large-scale load balancing, server consolidation, fault tolerance and other techniques. In a virtual cluster system, it's important to store the large number of VM images efficiently. The physical machines (host systems) and VMs (guest systems) may run with different OSes. You can have each VM installed on a remote server or replicated on multiple servers belonging to the same or different physical clusters. The boundary of a virtual cluster can change as you add, remove or dynamically migrate VM nodes over time.

PART-C

Q.15 What do you mean by virtualization of memory and I/O devices? Explain in details. [K.T.U. 2019]

Ans. Memory Virtualization : Virtual memory virtualization is similar to the virtual memory support provided by modern operating systems. In a traditional execution environment, the operating system maintains mappings of virtual memory to machine memory using page tables, which is a one-stage mapping from virtual memory to machine memory. All modern x86 CPUs include a memory management unit (MMU) and a translation lookaside buffer (TLB) to optimize virtual memory performance. However, in a virtual execution environment, virtual memory virtualization involves sharing the physical system memory in RAM and dynamically allocating it to the physical memory of the VMs.

That means a two-stage mapping process should be maintained by the guest OS and the VMM, respectively: virtual memory to physical memory and physical memory to machine memory. Furthermore, MMU virtualization should be supported, which is transparent to the guest OS. The guest OS continues to control the mapping of virtual addresses to the physical memory addresses of VMs. But the VMM is responsible for mapping the actual machine physical memory to the actual machine memory. Fig. (1) shows the two-level memory mapping procedure.




Fig. 1 : Two-level memory mapping procedure

Since each page table of the guest OSes has a separate page table in the VMM corresponding to it, the VMM page table is called the shadow page table. Nested page tables add another layer of indirection to virtual memory. The MMU already handles virtual-to-physical translations as defined by the OS. Then the physical memory addresses are translated to machine addresses using another set of page tables defined by the hypervisor. Since modern operating systems maintain a set of page tables for every process, the shadow page tables will get flooded. Consequently, the performance overhead and cost of memory will be very high.

VMware uses shadow page tables to perform virtual-memory-to-machine-memory address translation. Processors use TLB hardware to map the virtual memory directly to the machine memory to avoid the two levels of translation on every access. When the guest OS changes the virtual memory to a physical memory mapping, the VMM updates the shadow page tables to enable a direct lookup. The AMD Barcelona processor has featured hardware-assisted memory virtualization since 2007. It provides hardware assistance to the two-stage address translation in a virtual execution environment by using a technology called nested paging.

Example : Extended Page Table by Intel for Memory Virtualization : Since the efficiency of the software shadow page table technique was too low, Intel developed a hardware-based EPT technique to improve it, as illustrated in Figure. In addition, Intel offers a Virtual Processor ID (VPID) to improve use of the TLB. Therefore, the performance of memory virtualization is greatly improved. In Fig. (2), the page tables of the guest OS and EPT are all four-level.

When a virtual address needs to be translated, the CPU will first look for the L4 page table pointed to by Guest CR3. Since the address in Guest CR3 is a physical address in the guest OS, the CPU needs to convert the

Guest CR3 GPA to the host physical address (HPA) using EPT. In this procedure, the CPU will check the EPT TLB to see if the translation is there. If there is no required translation in the EPT TLB, the CPU will look for it in the EPT. If the CPU cannot find the translation in the EPT, an EPT violation exception will be raised.




Fig. 2 : Two-level memory mapping procedure

When the GPA of the L4 page table is obtained, the CPU will calculate the GPA of the L3 page table by using the GVA and the content of the L4 page table. If the entry corresponding to the GPA in the L4 page table is a page fault, the CPU will generate a page fault interrupt and will let the guest OS kernel handle the interrupt. When the PGA of the L3 page table is obtained, the CPU will look for the EPT to get the HPA of the L3 page table, as described earlier. To get the HPA corresponding to a GVA, the CPU needs to look for the EPT five times, and each time, the memory needs to be accessed four times. Therefore, there are 20 memory accesses in the worst case, which is still very slow. To overcome this shortcoming, Intel increased the size of the EPT TLB to decrease the number of memory accesses.

I/O Virtualization : I/O virtualization involves managing the routing of I/O requests between virtual devices and the shared physical hardware. At the time of this writing, there are three ways to implement I/O virtualization: full device emulation, para-virtualization, and direct I/O. Full device emulation is the first approach for I/O virtualization. Generally, this approach emulates well-known as real-world devices.

All the functions of a device or bus infrastructure, such as device enumeration, identification, interrupts, and DMA, are replicated in software. This software is located in the VMM and acts as a virtual device. The I/O access requests of the guest OS are trapped in the VMM which

interacts with the I/O devices. The full device emulation approach is shown in Fig. (3).




Fig. 3 : Device emulation for I/O virtualization implemented inside the middle layer that maps real I/O devices into the virtual devices for the guest device driver to use.

A single hardware device can be shared by multiple VMs that run concurrently. However, software emulation runs much slower than the hardware it emulates. The para-virtualization method of I/O virtualization is typically used in Xen. It is also known as the split driver model consisting of a frontend driver and a backend driver. The frontend driver is running in Domain 0. They interact with each other via a block of shared memory. The frontend driver manages the I/O requests of the guest OSes and the backend driver is responsible for managing the real I/O devices and multiplexing the I/O data of different VMs. Although para-I/O-virtualization achieves better device performance than full device emulation, it comes with a higher CPU overhead.

Direct I/O virtualization lets the VM access devices directly. It can achieve close-to-native performance without high CPU costs. However, current direct I/O virtualization implementations focus on networking for mainframes. There are a lot of challenges for commodity hardware devices. For example, when a physical device is reclaimed (required by workload migration) for later reassignment, it may have been set to an arbitrary state (e.g., DMA to some arbitrary memory locations) that can function incorrectly or even crash the whole system. Since software-based I/O virtualization requires a very high overhead of device emulation, hardware-assisted I/O virtualization is critical. Intel VT-d supports the remapping of I/O DMA transfers and device-generated interrupts. The architecture of VT-d provides the flexibility to support multiple usage models that may run unmodified, special-purpose, or "virtualization-aware" guest OSes.

Another way to help I/O virtualization is via self-virtualized I/O (SV-IO). The key idea of SV-IO is to harness the rich resources of a multicore processor. All tasks associated with virtualizing an I/O device are encapsulated in SV-IO. It provides virtual devices and an associated access API to VMs and a management API to the VMM. SV-IO defines one virtual interface (VIF) for every kind of virtualized I/O device, such as virtual network interfaces, virtual block devices (disk), virtual camera devices, and others. The guest OS interacts with the VIFs via VIF device drivers. Each VIF consists of two message queues. One is for outgoing messages to the devices and the other is for incoming messages from the devices. In addition, each VIF has a unique ID for identifying it in SV-IO.

Example : VMware Workstation for I/O Virtualization : The VMware Workstation runs as an application. It leverages the I/O device support in guest OSes, host OSes, and VMM to implement I/O virtualization. The application portion (VMAp) uses a driver loaded into the host operating system (VMDriver) to establish the privileged VMM, which runs directly on the hardware. A given physical processor is executed in either the host world or the VMM world, with the VMDriver facilitating the transfer of control between the two worlds. The VMware Workstation employs full device emulation to implement I/O virtualization. Fig. (4) shows the functional blocks used in sending and receiving packets via the emulated virtual NIC.

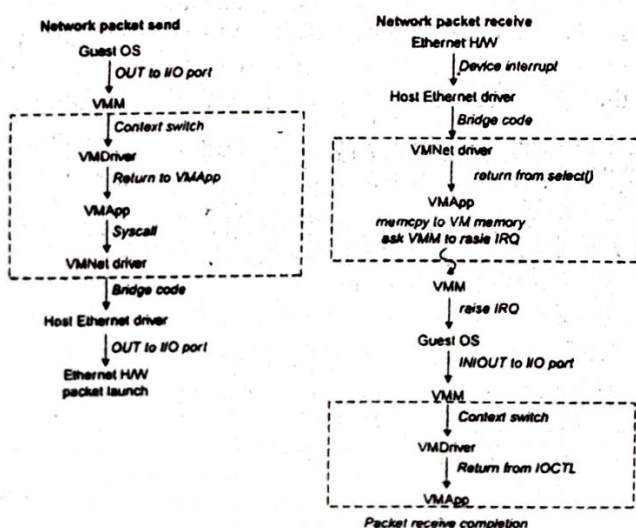


Fig. 4 : Functional blocks involved in sending and receiving network packets

The virtual NIC models an AMD Lance Am79C970A controller. The device driver for a Lance controller in the guest OS initiates packet transmissions by reading and writing a sequence of virtual I/O ports; each read or write switches back to the VMAp to emulate the Lance port accesses. When the last OUT instruction of the sequence is encountered, the Lance emulator calls a normal write() to the VMNet driver. The VMNet driver then passes the packet onto the network via a host NIC and then the VMAp switches back to the VMM. The switch raises a virtual interrupt to notify the guest device driver that the packet was sent. Packet receives occur in reverse.

Q.16 What do you mean by hypervisor? Explain different types of hypervisors. [R.T.U. 2019, 2017]

Ans. Hypervisor Architecture : The hypervisor supports hardware-level virtualization on bare metal devices like CPU, memory, disk and network interfaces. The hypervisor software sits directly between the physical hardware and its OS. This virtualization layer is referred to as either the VMM or the hypervisor. The hypervisor provides hypercalls for the guest OSes and applications. Depending on the functionality, a hypervisor can assume a micro-kernel architecture like the Microsoft Hyper-V. Or it can assume a monolithic hypervisor architecture like the VMware ESX for server virtualization. A micro-kernel hypervisor includes only the basic and unchanging functions (such as physical memory management and processor scheduling). The device drivers and other changeable components are outside the hypervisor. A monolithic hypervisor implements all the aforementioned functions, including those of the device drivers. Therefore, the size of the hypervisor code of a micro-kernel hypervisor is smaller than that of a monolithic hypervisor. Essentially, a hypervisor must be able to convert physical devices into virtual resources dedicated for the deployed VM to use.

Types of Hypervisor :

(i) VMWare Hypervisor: VMware ESXi is an operating system independent hypervisor based on the VMkernel operating system interfacing with agents that run a top it. ESXi is the exclusive hypervisor for VMware vSphere 5.x licenses. VMware describes an ESXi system as similar to a stateless compute node. State information can be uploaded from a saved configuration file. ESXi's VMkernel interfaces directly with VMware agents and approved third party modules. Virtualization administrators can configure VMware ESXi through its console or the VMware vSphere Client and check VMware's Hardware

environment, the VM state is akin to a tree: At any point, execution can go into N different branches where multiple instances of a VM can exist at any point in this tree at any given time. VMs are allowed to roll back to previous states in their execution (e.g., to fix configuration errors) or rerun from the same point many times (e.g., as a means of distributing dynamic content or circulating a "live" system image).

Q.18 What are the different types of virtualization?

[R.T.U. 2017]

Ans. Different Types of Virtualization

- 1. Server Virtualization :** This type is where most of the attention is focused right now in the world of virtualization and is where most companies begin an implementation of this technology. That's not very shocking in light of the fact that server sprawl has become a very large and legitimate problem in enterprises throughout the world. Where a company is simply running out of room in which to place all of their servers, this type of virtualization would of course get viewed with strong interest. Because each server typically serves one function (i.e., mail server, file server, Internet server, enterprise resource planning server, etc.), with each server using only a fraction of its true processing power, server virtualization breaks through the "one application, one server" barrier and facilitates the consolidation of numerous servers into one physical server. This equates to (a) less physical servers required and (b) 70 to 80 percent or higher utilization of existing hardware as opposed to the previous 10 to 15 percent.
- 2. Application Virtualization :** An application runs on another host from where it is installed in a variety of ways. It could be done by application streaming, desktop virtualization or VDI or a VM package (like VMware ACE creates with a player). Microsoft Softgrid is an example of Application virtualization.
- 3. Presentation Virtualization :** This is what Citrix Met frame (and the ICA protocol) as well as Microsoft Terminal Services (and RDP) are able to create. With presentation virtualization, an application actually runs on another host and all that you see on the client is the screen from where it is run.
- 4. Network Virtualization :** With network virtualization, the network is "carved up" and can be used for multiple purposes such as running a protocol analyzer inside an Ethernet switch.

Components of a virtual network could include NICs, switches, VLANs, network storage devices, virtual network containers and network media.

- 5. Storage Virtualization :** With storage virtualization, the disk/data storage for your data is consolidated to and managed by a virtual storage system. The servers connected to the storage system aren't aware of where the data really is. Storage virtualization is sometimes described as "abstracting the logical storage from the physical storage."

With virtualization, the virtual machine uses hardware directly, although there is an overarching scheduler. As such, no emulation is taking place, but this limits what can be run inside virtual machines to operating systems that could otherwise run atop the underlying hardware. That said, this method provides the best overall performance of the two solutions.

With emulation, since an entire machine can be created as a virtual construct, there is a wider variety of opportunities, but with the aforementioned 'emulation penalty. But, emulation makes it possible to, for example, run programs designed for a completely different architecture on an x86 PC. This approach is common, for example, when it comes to running old games designed for obsolete platforms on today's modern systems. Because everything is emulated in software, there is a performance hit in this method, although today's massively powered processors often cover for this.

Both methods are used for various purposes and are sometimes confused, so be aware of the differences.

There are two types of hypervisors: Type 1 and Type 2.

Type 1 hypervisors run directly on the system hardware. They are often referred to as a "native" or "bare metal" or "embedded" hypervisors in vendor literature.

Type 2 hypervisors run on a host operating system. When the virtualization movement first began to take off, Type 2 hypervisors were most popular. Administrators could buy the software and install it on a server they already had.

Type 1 hypervisors are gaining popularity because building the hypervisor into the firmware is proving to be more efficient. According to IBM, Type 1 hypervisors provide higher performance, availability, and security than Type 2 hypervisors. (IBM recommends that Type 2 hypervisors be used mainly on client systems where efficiency is less critical or on systems where support for a broad range of I/O devices is important and can be provided by the host operating system.)

Experts predict that shipping hypervisors on bare metal will impact how organizations purchase servers in the future. Instead of selecting an OS, they will simply

have to order a server with an embedded hypervisor and run whatever OS they want.

To keep their market-share, each of the major virtualization software vendors has announced plans to work with hardware manufacturers to embed their hypervisor into the manufacturer's firmware.

Q.19 What is virtualization? Explain the various implementation levels of virtualization. Differentiate between server and desktop virtualization. [R.T.U. 2015]

Ans. Virtualization : Refer to Q.13.

Levels of Virtualization Implementation : A traditional computer runs with a host operating system specially tailored for its hardware architecture. After virtualization, different user applications managed by their own operating systems (guest OS) can run on the same hardware, independent of the host OS. This is often done by adding additional software, called a virtualization layer. This virtualization layer is known as hypervisor or virtual machine monitor (VMM). The VMs are shown in the upper boxes, where applications run with their own guest OS over the virtualized CPU, memory and I/O resources. The main function of the software layer for virtualization is to virtualize the physical hardware of a host machine into virtual resources to be used by the VMs, exclusively. This can be implemented at various operational levels, as we will discuss shortly. The virtualization software creates the abstraction of VMs by interposing a virtualization layer at various levels of a computer system. Common virtualization layers include the instruction set architecture (ISA) level, hardware level, operating system level, library support level and application level (see Figure).

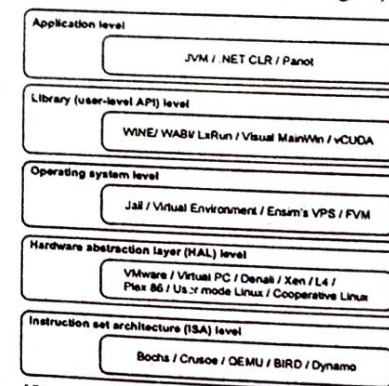


Fig. : Virtualization ranging from hardware to applications in five abstraction levels

Q.21 Explain HDFS in detail.

Ans. The Hadoop Distributed File System (HDFS) is a distributed file system designed to run on commodity hardware. It has many similarities with existing distributed file systems. However, the differences from other distributed file systems are significant. HDFS is highly fault-tolerant and is designed to be deployed on low-cost hardware. HDFS provides high throughput access to application data and is suitable for applications that have large data sets. HDFS relaxes a few POSIX requirements to enable streaming access to file system data. HDFS was originally built as infrastructure for the Apache Nutch web search engine project. HDFS is now an Apache Hadoop.

HDFS is a Java-based file system that provides scalable and reliable data storage and it was designed to span large clusters of commodity servers. HDFS has demonstrated production scalability of up to 200 PB of storage and a single cluster of 4500 servers, supporting close to a billion files and blocks. When that quantity and quality of enterprise data is available in HDFS and YARN enables multiple data access applications to process it, Hadoop users can confidently answer questions that eluded previous data platforms.

In HDP 2.2, the rolling upgrade feature and the underlying HDFS High Availability configuration enable Hadoop operators to upgrade the cluster software and restart upgraded services, without taking the entire cluster down.

Goals of HDFS

Hardware Failure: Hardware failure is the norm rather than the exception. An HDFS instance may consist of hundreds or thousands of server machines, each storing part of the file system's data. The fact that there are a huge number of components and that each component has a non-trivial probability of failure means that some component of HDFS is always non-functional. Therefore, detection of faults and quick, automatic recovery from them is a core architectural goal of HDFS.

Streaming Data Access: Applications that run on HDFS need streaming access to their data sets. They are not general purpose applications that typically run on general purpose file systems. HDFS is designed more for batch processing rather than interactive use by users. The emphasis is on high throughput of data access rather than low latency of data access. POSIX imposes many hard requirements that are not needed for applications that are targeted for HDFS. POSIX semantics in a few key areas has been traded to increase data throughput rates.

Cloud Computing

Large Data Sets: Applications that run on HDFS have large data sets. A typical file in HDFS is gigabytes to terabytes in size. Thus, HDFS is tuned to support large files. It should provide high aggregate data bandwidth and scale to hundreds of nodes in a single cluster. It should support tens of millions of files in a single instance.

Simple Coherency Model: HDFS applications need a write-once-read-many access model for files. A file once created, written and closed need not be changed. This assumption simplifies data coherency issues and enables high throughput data access. A MapReduce application or a web crawler application fits perfectly with this model. There is a plan to support append-only writes to files in the future.

"Moving Computation is Cheaper than Moving Data"

A computation requested by an application is much more efficient if it is executed near the data it operates on. This is especially true when the size of the data set is huge. This minimizes network congestion and increases the overall throughput of the system. The assumption is that it is often better to migrate the computation closer to where the data is located rather than moving the data to where the application is running. HDFS provides interfaces for applications to move themselves closer to where the data is located.

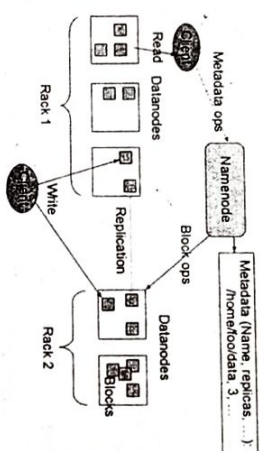
Portability Across Heterogeneous Hardware and Software Platforms

HDFS has been designed to be easily portable from one platform to another. This facilitates widespread adoption of HDFS as a platform of choice for a large set of applications.

NameNode and DataNodes: HDFS has a master/slave architecture. An HDFS cluster consists of a single NameNode, a master server that manages the file system namespace and regulates access to files by clients. In addition, there are a number of DataNodes, usually one per node in the cluster, which manage storage attached to the nodes that they run on. HDFS exposes a file system namespace and allows user data to be stored in files. Internally, a file is split into one or more blocks and these blocks are stored in a set of DataNodes. The NameNode executes file system namespace operations like opening, closing and renaming files and directories. It also determines the mapping of blocks to DataNodes. The DataNodes are responsible for serving read and write requests from the file system's clients. The DataNodes also perform block creation, deletion and replication upon instruction from the NameNode.

The NameNode and DataNode are pieces of software designed to run on commodity machines. These

machines typically run a GNU/Linux operating system (OS). HDFS is built using the Java language; any machine that supports Java can run the NameNode or the DataNode software. Usage of the highly portable Java language means that HDFS can be deployed on a wide range of machines. A typical deployment has a dedicated machine that runs only the NameNode software. Each of the other machines in the cluster runs one instance of the DataNode software. The architecture does not preclude running multiple DataNodes on the same machine but in a real deployment that is rarely the case.

HDFS Architecture**Fig.: HDFS architecture**

The existence of a single NameNode in a cluster greatly simplifies the architecture of the system. The NameNode is the arbitrator and repository for all HDFS metadata. The system is designed in such a way that user data never flows through the NameNode.

The File System Namespace: HDFS supports a traditional hierarchical file organization. A user or an application can create directories and store files inside these directories. The file system namespace hierarchy is similar to most other existing file systems; one can create and remove files, move a file from one directory to another or rename a file. HDFS does not yet implement user quotas. HDFS does not support hard links or soft links. However, the HDFS architecture does not preclude implementing these features.

The NameNode maintains the file system namespace. Any change to the file system namespace or its properties is recorded by the NameNode. An application can specify the number of replicas of a file that should be maintained by HDFS. The number of copies of a file is called the replication factor of that file. This information is stored by the NameNode.

Data Replication: HDFS is designed to reliably store very large files across machines in a large cluster. It stores each file as a sequence of blocks; all blocks in a file except the last block are the same size. The blocks of a file are

CLC-49

replicated for fault tolerance. The block size and replication factor are configurable per file. An application can specify the number of replicas of a file. The replication factor can be specified at file creation time and can be changed later. Files in HDFS are write-once and have strictly one writer at any time.

The NameNode makes all decisions regarding replication of blocks. It periodically receives a Heartbeat and a Blockreport from each of the DataNodes in the cluster. Receipt of a Heartbeat implies that the DataNode is functioning properly. A Blockreport contains a list of all blocks on a DataNode.

Q.22 What are the needs of virtualization? Define different types of virtualization.

Ans. Industry will continue to adopt virtualization for many reasons: collections of inefficient servers can be replaced with fewer machines; software can be tested while isolated in harmless virtual partitions; and data centers can gracefully (and virtually) conform for shifting work models, new technologies and changing corporate priorities.

The future of enterprise IT management will be based on virtual computing. Intel VT makes it possible to maximize computer utilization while minimizing all associated overheads of management, power consumption, maintenance and physical space.

Intel Virtualization Technology provides a comprehensive roadmap to address virtualization challenges and includes support for CPU and I/O virtualization and a strong VMM ecosystem. Intel was the first and is the leading provider of hardware support for virtualization technologies.

Today's IT intensive enterprise must always be on the lookout for the latest technologies that allow businesses to run with fewer resources while providing the infrastructure to meet today and future customer needs. Virtualization utilizing Intel Virtualization Technology is the cutting edge of enterprise information technology. Intel is closely working with VMware, XENSource, Jaluna, Parallel, tenAys, VirtualIron, RedHat, Novell and other VMM developers.

Needs of Virtualization

Server Consolidation: It is not unusual to achieve 10:1 virtual to physical machine consolidation. This means that ten server applications can be run on a single machine that had required as many physical computers to provide the unique operating system and technical specification environments in order to operate. Server utilization is optimized and legacy software can maintain old OS configurations while new applications are running in VMs with updated platforms.

SECURITY IN CLOUD

4

YEARS QUESTIONS

- Logging
- IP restrictions
- API gateways
- CASB

Q4 What do you mean by SLA?

Ans. A Service Level Agreement (SLA) is a contract between a service provider and its internal or external customers that documents what services the provider will furnish.

PART-B

Q5 Explain the cloud security requirements and give fundamental model for cloud information security. [R.T.U. 2019]

OR
What are cloud security requirements? Explain the various security challenges in cloud computing. [R.T.U. 2015]

OR
Explain cloud information security/fundamentals introduced in cloud security management.

Ans. Cloud security requirements : Cloud computing opens up a new world of opportunities for businesses, but mixed in with these opportunities are numerous security requirements (and the challenges associated with them) that need to be considered and addressed prior to commit to a cloud computing strategy. Cloud computing security requirements/challenges fall into three broad categories:

Cloud Computing

1. **Data Protection:** Securing your data both at rest and in transit
2. **User Authentication:** Limiting access to data and monitoring who accesses the data
3. **Disaster and Data Breach:** Contingency Planning

1. Data Protection

Implementing a cloud computing strategy means placing critical data in the hands of a third party, so ensuring the data remains secure both at rest (data residing on storage media) as well as when in transit is of paramount importance. Data needs to be encrypted at all times, with clearly defined roles when it comes to who will be managing the encryption keys. In most cases, the only way to truly ensure confidentiality of encrypted data that resides on a cloud provider's storage servers is for the client to own and manage the data encryption keys.

2. User Authentication

Data resting in the cloud needs to be accessible only by those authorized to do so, making it critical to both restrict and monitor who will be accessing the company's data through the cloud. In order to ensure the integrity of user authentication, companies need to be able to view data access logs and audit trails to verify that only authorized users are accessing the data. These access logs and audit trails additionally need to be secured and maintained for as long as the company needs or legal purposes require. As with all cloud computing security challenges, it's the responsibility of the customer to ensure that the cloud provider has taken all necessary security measures to protect the customer's data and the access to that data.

3. Disaster and Data Breach

With the cloud serving as a single, centralized repository for a company's mission critical data, the risks of having that data compromised due to a data breach or temporarily made unavailable due to a natural disaster are real concerns. Much of the liability for the disruption of data in a cloud ultimately rests with the company whose mission critical operations depend on that data, although liability can and should be negotiated in a contract with the services provider prior to commitment. A comprehensive security assessment from a neutral third party is strongly recommended as well.

Companies need to know how their data is being secured and what measures the service provider will be taking to ensure the integrity and availability of that data should the unexpected occur. Additionally, companies should also have contingency plans in place in the event their cloud provider fails or goes bankrupt. Can the data be easily retrieved and migrated to a new service provider or to a non cloud strategy if this happens? And what

CLC-57

happens to the data and the ability to access that data if the provider gets acquired by another company?

Cloud Information Security Objectives

Developing secure software is based on applying the secure software design principles that form the fundamental basis for software assurance. Software assurance has been given many definitions and it is important to understand the concept. The Software Security Assurance Report2 defines software assurance as "the basis for gaining justifiable confidence that software will consistently exhibit all properties required to ensure that the software, in operation, will continue to operate dependably despite the presence of sponsored (intentional) faults. In practical terms, such software must be able to resist most attacks, tolerate as many as possible of those attacks it cannot resist and contain the damage and recover to a normal level of operation as soon as possible after any attacks it is unable to resist or tolerate." The U.S. Department of Defense (DoD) Software Assurance Initiative3 defines software assurance as "the level of confidence that software functions as intended and is free of vulnerabilities, either intentionally or unintentionally designed or inserted as part of the software." The Data and Analysis Center for Software (DACS) requires that software must exhibit the following three properties to be considered secure:

Dependability : Software that executes predictably and operates correctly under a variety of conditions, including when under attack or running on a malicious host.

Trustworthiness : Software that contains a minimum number of vulnerabilities or no vulnerabilities or weaknesses that could sabotage the software's dependability. It must also be resistant to malicious.

Survivability (Resilience) : Software that is resistant to or tolerant of attacks and has the ability to recover as quickly as possible with as little harm as possible seven complementary principles that support information assurance are confidentiality, integrity, availability, authentication, authorization, auditing and accountability. These concepts are summarized in the following sections.

Confidentiality, Integrity and Availability

Confidentiality, integrity and availability are sometimes known as the CIA triad of information system security and are important pillars of cloud software assurance. Confidentiality refers to the prevention of intentional or unintentional unauthorized disclosure of information. Confidentiality in cloud systems is related to the areas of intellectual property rights, covert channels, traffic analysis, encryption and inference.

Intellectual Property Rights : Intellectual property (IP) includes inventions, designs, artistic, musical and literary

CLC-58

works. Rights to intellectual property are covered by copyright laws, which protect creations of the mind and patents, which are granted for new inventions.

Covert channels : A covert channel is an unauthorized and unintended communication path that enables the exchange of information. Covert channels can be accomplished through timing of messages or inappropriate use of storage mechanisms.

Traffic analysis : Traffic analysis is a form of confidentiality breach that can be accomplished by analyzing the volume, rate, source and destination of message traffic, even if it is encrypted. Increased message activity and high bursts of traffic can indicate a major event is occurring. Counter measures to traffic analysis include maintaining a near-constant rate of message traffic and disguising the source and destination locations of the traffic.

Encryption : Encryption involves scrambling messages so that they cannot be read by an unauthorized entity, even if they are intercepted. The amount of effort (work factor) required to decrypt the message is a function of the strength of the encryption key and the robustness and quality of the encryption algorithm.

Inference : Inference is usually associated with database security. Inference is the ability of an entity to use and correlate information protected at one level of security to uncover information that is protected at a higher security level.

Q6 Explain different legal issues in cloud computing. [R.T.U. 2019]

Ans. Legal Issues for the Cloud

Service levels : It should go without saying that the starting point should be the business case and intended use of the service, and not any legal document, such as a service level agreement (SLA). Understand what business problem the service will be solving; the intended internal and external users; when, where and how the service will be accessed; whether or not the service is business-critical; the practical consequences if the service is down or degraded for any period of time; and how the use of the service may change over time. Then, ensure the SLA reflects your needs.

Termination or suspension of service : The software application and/or the data running or housed in the cloud may be critical to your business. Continuity of access and use (to both the application and data), especially when both are on a third-party server, are of utmost importance. To that end, does the cloud vendor in each instance notify

you when any of the terms of the agreement may have been violated, and are you given an opportunity to remedy each violation?

Representations and warranties; indemnities : While seemingly arcane, in terms of potential pitfalls, these provisions may be the most important. A representation is a statement of fact, either past or present, while a warranty may express a promise. Typical reps and warranties should confirm that there are no pending or threatened claims of intellectual property right (IPR) infringement (after all, who wants legal problems on day one?) and address continued no infringement, performance (as to the underlying app), and data security and privacy.

Breach of a warranty will typically give rise to a limited remedy and thus will be to the exclusion of other remedies, such as money damages. Therefore, be sure the limited remedy makes business sense and will suffice. Note also that cloud providers typically request reps and warranties from the customer, including those pertaining to the customer's data. To that end, the buyer must be careful about the sources of its data or risk exposing itself to liability.

An indemnity is a contractual obligation to compensate a party for a loss. Thus, an indemnity would compensate the cloud customer for any claims that its use of the service violated any third-party IP rights, such as patent, copyright or trademark. These suits (especially patent) are costly, so care must be taken to ensure that you are adequately covered.

Confidentiality : Cloud customers should be sure to get satisfactory promises regarding which vendor personnel will have access to confidential information (including customer data) and what steps the vendor will undertake to maintain the confidentiality of that information. Data is king, and this provision deserves considerable attention.

Commercial/other : The considerations above are a good starting point but they are just the tip of the iceberg. Here are a few more to consider: storage fees, if and when there are automatic upgrades; whether or not there are multiple environments (e.g., development, test, and production) available to customer, how customization works in a cloud setting, how many data recoveries does the vendor provide free of charge (and what are the costs of additional backups), and how easy is it to move to another cloud and how will the vendor support the transition?

Q.7 Explain business continuity planning and Disaster recovery planning. [R.T.U. 2019]

OR

What is business continuity planning (BCP)? Explain the importance and process of BCP? [R.T.U. 2015]

OR

Write short notes on BCP. [R.T.U. 2018]

Ans. Business Continuity Planning is the act of proactively working out a way to prevent, if possible and manage the consequences of a disaster, limiting it to the extent that a business can afford.

BCP Process

The first step is to identify the assets and processes that are critical to the business, some may have been identified during recent risk assessment exercises. The BCP should answer the following questions.

1. Which roles and individuals are vital for fulfilling business commitments?
2. What equipment, IT, transport etc will staff need to maintain operations?
3. How long can the business function before full operations are restored?
4. What alternative resources are available?
5. Which departments are vital for fulfilling orders and contractual obligations?
6. Which suppliers and other third parties are integral to daily routines?

The BCP should then present all the steps that staff are expected to follow in the aftermath of an incident in order to maintain essential operations and return to 'business as usual' as soon as possible.

The structure and detail of the BCP will vary from business to business and from location to location but, in general, should include:

The initial response

1. Clear roles and delegated responsibilities for those who will take charge of co-ordinating the initial response and from where in each location.
2. Do those working with emergency services have access to any prepared response packs?
3. Factors determining part or full evacuation of premises.
4. How first-aid will be provided?
5. Arranging internal and external communications.
6. Arrangements for marshalling crowds towards pre-determined muster points.
7. Support for people with disabilities, restricted mobility or other needs.

Longer-term planning

1. When and how alternative accommodation and facilities will be utilised?

2. Arrangements for maintaining access to key records and IT systems.
3. Contingency arrangements for critical operations such as financial transactions, client orders, receipt of deliveries, production commitments.
4. How the appropriate staff will be able to access any contingency sites?
5. Agreed procedures for re-commencing routine operations.
6. How staff kept off-site from site can be kept informed?

Importance of BCP

There are various threats and vulnerabilities to which business today is exposed. They could be:

1. catastrophic events such as floods, earthquakes or acts of terrorism
2. accidents or sabotage
3. outages due to an application error, hardware or network failures

Some of them come unwarned. Most of them never happen. The key is to be prepared and be able to respond to the event when it does happen, so that the organization survives; its losses are minimized; it remains viable and it can be "business as usual", even before the customers feel the effects of the downtime. An effective Business Continuity Plan serves to secure businesses against financial disasters.

The bonus – Customer satisfaction, enhanced corporate image and no dip in the market share.

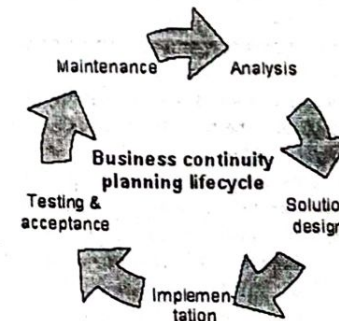


Fig.

Cloud Disaster Recovery Approaches : While all cloud disaster recovery plans involve the cloud in some way, the specific types of cloud resources that are used, and the manner in which they are deployed, varies depending on the DR approach you take.

There are four approaches to consider:

Backup and recovery : This is the most straightforward offsite disaster recovery strategy. It involves backing up data to the cloud and recovering it from the cloud when a disaster occurs. (As noted above, keep in mind that the data could be recovered from your backup location to a local on-premise environment, or to another cloud infrastructure, depending on your needs.) For a simple backup and recovery strategy to work, you need to ensure that it meets your RTO and RPO requirements. In cases where RPO and RTO needs are very high, a backup and recovery approach to offsite disaster recovery may be difficult to achieve.

Pilot light : Under this approach, you keep a copy of your production virtual servers and databases stored in the cloud at all times. You keep their data and configurations synced with those of the production systems, but the cloud-based backup resources are spun up and activated only in the event that your production systems fail due to a disaster. Thus, the cloud backup environment functions like a pilot light in your furnace. It is always ready to fire up on demand; however, because it will take some time to spin up the cloud-based backup resources, recovery following a disaster is not instantaneous. The trade-off for the delay is that you don't need to pay to have your backup resources running in the cloud constantly.

Warm standby : Warm standby is similar to the pilot light approach, except that your backup virtual servers and databases are actually running at all times. This enables you to put the backup resources into production almost instantaneously whenever a disaster strikes. However, this approach will be more expensive, because you have to pay to keep backup copies of your infrastructure running in the cloud at all times.

Multi-site : This approach entails using the warm standby technique, but instead of having only one copy of your workloads running in the cloud at all times, you run multiple copies that are spread across different geographic zones of the cloud. This strategy provides the greatest level of availability because it ensures that you can restore your workloads very quickly even if part of your cloud provider's infrastructure has failed. It is also, however, the most expensive approach.

Ans. Threats in Cloud : Mo level of any implementation security concerns of that system or any hardware, along with se whole new level of alarming th of obstacles to a company's gro helped many companies scale the security vulnerabilities.

Data Breach : The year 2018 wi year of the maximum data organizations like Facebook was which involved breaching of milli This data breach shocked the en people started questioning the safe of these big corporations. One mor reported by Equifax where it releas that in 2017, around 143 million compromised. This is the problem yo to maintain a humongous database than 100 million people.

Trojan Horse : The main pillar for human resources, that is, its emplo work in a protocol that helps the comp functioning. But, at times they are p wrong route and might leak their co confidential data through various mea or external pressure. This cannot be a won't be able to detect and tackle i company to take some extreme steps data with extra precaution.

Vulnerable APIs : Cloud Service Pr manage a collection of application progra (APIs) that customers use to operate a with the cloud services. Companies use orchestrate, and keep a track of their ass users. These APIs can be accessed thro by their CSPs. They are equally vulnerabl system and its libraries which can be anywhere online.

Support Issues : Companies can fac compatibility when launching their applicati

There could be certain products t need to launch within a span of time. Imp same in the cloud could be feasible but it the company some time. For launching it i should be able to support the cloud in mea orchestrators, functions and its level of con **Insufficient Data Backup :** Data in the cl in many forms and at different locations. T located anywhere and can be accessed from

- (2) **Work with a Third Party to assure Cloud Security on a Regular Basis** : By having multiple parties, it increases the security risks. However, small and medium businesses without large IT departments, sometimes need assistance to audit and ensure cloud security. For some industries, this assistance comes in the form of industry-standard security certification. You should utilize third-party audits to ensure that your cloud provider is following your industry's standards of security.
- (3) **Implement end-to-end Encryption** : The end-to-end encryption is particularly for cloud storage. It decreases the likelihood of our data being breached. Most cloud storage solutions have encrypted data upload and downloads, but does not store the data encrypted. The method with the least amount of risk requires our data to be encrypted prior to upload, while it is in provider's datacenter, and is only decrypted with a required encryption key.
- (4) **Regularly Update In-House Software** : If we are running outdated operating systems such as Windows XP and outdated internet browsers such as IE 7, we could be at risk despite encryption and third party audits.

PART-C

Q.13 Describe all the possible attacks that can be used to disrupt the cloud. [R.T.U. 2019]

OR

What are the types of security policies for cloud computing? [R.T.U. 2015]

OR

Explain Security Policies for cloud computing. [R.T.U. 2018]

Ans. Cloud computing opens up a new world of opportunities for businesses, but mixed in with these opportunities are numerous security challenges that need to be considered and addressed prior to commit a cloud computing strategy. While there are real benefits of using cloud computing, including some key security advantages, there are just as many if not more security challenges that prevent customers from committing to a cloud computing strategy. Ensuring that your data is securely protected both at rest and in transit, restricting and monitoring access to that data via user authentication and access logging and adequately planning for the very real possibilities of compromised or inaccessible data due to

data breaches or natural disasters are all key security challenges that a company must address when considering cloud computing providers. Cloud computing security challenges as follows:

1. Data Breaches

The data breach at Target, resulting in the loss of personal and credit card information of up to 110 million individuals, was one of the series of startling thefts that took place during the normal processing and storage of data. "Cloud computing introduces significant new avenues of attack," said the CSA report authors. The absolute security of hypervisor operation and virtual machine operations is still to be proved. Indeed, critics question whether such absolute security can exist. The report's writers said there's lab evidence though none known in the wild that breaches via hypervisors and virtual machines may occur eventually.

2. Data Loss

A data breach is the result of a malicious and probably intrusive action. Data loss may occur when a disk drive dies without its owner having created a backup. It occurs when the owner of encrypted data loses the key that unlocks it. Small amounts of data were lost for some Amazon Web Service customers as its EC2 cloud suffered "a remirroring storm" due to human operator error on Easter weekend in 2011. And a data loss could occur intentionally in the event of a malicious attack.

3. Account or Service Traffic Hijacking

Account hijacking sounds too elementary to be a concern in the cloud, but CSA says it is a problem. Phishing, exploitation of software vulnerabilities such as buffer overflow attacks and loss of passwords and credentials can all lead to the loss of control over a user account. An intruder with control over a user account can eavesdrop on transactions, manipulate data, provide false and business-damaging responses to customers and redirect customers to a competitor's site or inappropriate sites.

4. Insecure APIs

The cloud era has brought about the contradiction of trying to make services available to millions while limiting any damage all these largely anonymous users might do to the service. The answer has been a public facing application programming interface or API, that defines how a third party connects an application to the service and providing verification that the third party producing the application is who he says he is.

Leading web developers, including ones from Twitter and Google, collaborated on specifying OAuth, an open authorization service for web services that controls third party access.

5. Denial of Service

Denial of service attacks are an old disrupter of online operations, but they remain a threat nevertheless. The assault by hundreds or thousands or millions of automated requests for service has to be detected and screened out before it ties up operations, but attackers have improvised increasingly sophisticated and distributed ways of conducting the assault, making it harder to detect which parts of the incoming traffic are the bad actors versus legitimate users.

6. Malicious Insiders

With the Edward Snowden case and NSA revelations in the headlines, malicious insiders might seem to be a common threat. If one exists inside a large cloud organization, the hazards are magnified. One tactic cloud customers should use to protect themselves is to keep their encryption keys on their own premises, not in the cloud.

"If the keys are not kept with the customer and are only available at data-usage time, the system is still vulnerable to malicious insider attack." Systems that depend "solely on the cloud service provider for security are at great risk" from a malicious insider, the report said.

7. Abuse of Cloud Services

Cloud computing brings large-scale, elastic services to enterprise users and hackers alike. "It might take an attacker years to crack an encryption key using his own limited hardware. But using an array of cloud servers, he might be able to crack it in minutes," the report noted. Or hackers might use cloud servers to serve malware, launch DDoS attacks or distribute pirated software.

8. Insufficient due Diligence

"Too many enterprises jump into the cloud without understanding the full scope of the undertaking," said the report. Without an understanding of the service providers' environment and protections, customers don't know what to expect in the way of incident response, encryption use and security monitoring. Not knowing these factors means "organizations are taking an unknown levels of risk in ways they may not even comprehend, but that are a far departure from their current risks," wrote the authors.

Chances are, expectations will be mismatched between customer and service. Enterprises may push applications that have internal on-premises network security controls into the cloud, where those network security controls don't work. If enterprise architects don't understand the cloud environment, their application designs may not function with proper security when they're run in a cloud setting, the report warned.

Multi-Paxos where no contention eliminates prepare phases and means the normal operational case is a steady stream of accept messages from the coordinator.

Chubby is a distributed lock service intended for coarse-grained synchronization of activities within Google's distributed systems.

Chubby has become Google's primary internal name service; it is a common rendezvous mechanism for systems such as MapReduce, the storage systems GFS and Bigtable use Chubby to elect a primary from redundant replicas; and it is a standard repository for files that require high availability, such as access control lists.

Chubby is a relatively heavy-weight system intended for coarse-grained locks, locks held for "hours or days", not "seconds or less."

The paper talks about many of the practical issues they encountered building this large-scale system. It is a good read.

Developers sometimes do not plan for high availability in the way one would wish. Often their systems start as prototypes with little load and loose availability guarantees; invariably the code has not been specially structured for use with a consensus protocol. As the service matures and gains clients, availability becomes more important; replication and primary election are then added to an existing design.

Developers are often unable to predict how their services will be used in the future and how use will grow. A module written by one team may be reused a year later by another team with disastrous results.

Despite attempts at education, developers regularly write loops that retry indefinitely when a file is not present or poll a file by opening it and closing it repeatedly when one might expect they would open the file just once.

Developers rarely consider availability. We find that our developers rarely think about failure probabilities.

Developers also fail to appreciate the difference between a service being up and that service being available to their applications.

Unfortunately, many developers choose to crash their applications on receiving [a failover] event, thus decreasing the availability of their systems substantially.

Q.20 Explain QoS (Quality of Services) monitoring in cloud.

Ans. The evolution of Cloud Computing environments has resulted in a new impulse to the service oriented computing, with hardware resources, whole applications and entire business processes provided as services in the so called "as a service" paradigm. In such a paradigm the resulting

interactions should involve actors (users and providers of services) belonging to different entities and possibly to different companies, hence the success of such a new vision of the IT world is strictly tied to the possibility of guaranteed high quality levels in the provisioning of resources and services.

Quality of Service (QoS) monitoring is key for a company's success, since assessing the actual quality of what service users are paying for has become a mission-critical business practice requirement and it will be even more so in the future. The ever increasing complexity of individual components and interconnection among them has impaired our ability to measure the Key Performance Indicators (KPIs) of the service which is to be monitored. Emerging development paradigms, together with the amazing increase of the scale, have made this task even more challenging in upcoming Future Internet (FI) scenarios, since individual business processes - whose internals are completely unknown - are being integrated to quickly implement and cheaply deploy semantically richer business processes.

Top Management Challenges

Inability to identify applications that could be seamlessly moved to the cloud : Before making decisions about applications that should be moved to the cloud environment, organizations should make a calculation about IT and business benefits that they can achieve from this action. Additionally, organizations should have capabilities in place to test whether the cloud infrastructure they are using can support applications that are being transferred to the cloud.

Inability to make educated decisions about adding or terminating cloud resources : Deploying cloud computing services changes the way organizations go about managing their computing resources, as it gives them more flexibility in using available capacity in the way that is the most cost effective. Instead of making costly investments in new hardware when they need additional capacity, organizations have the ability to increase and decrease cloud resources used as the demand changes.

Inability to monitor performance of applications that use a hybrid cloud approach : Organizations using cloud computing services need to have visibility not only into the performance of applications that have moved to the cloud, but also into the different computing resources on which these applications depend. Typically, organizations find it easier to monitor the performance of applications that are hosted at a single server as opposed to the performance of composite applications that are pulling computing resources from different sources. This issue becomes even more complex if computing resources are

hosted outside of corporate firewalls and organizations do not have a full control and visibility into the performance of these applications.

Organizations sometimes use a hybrid model for deploying cloud computing, which presents end-user organizations with the challenge of monitoring usage of resources that are hosted and managed both externally and internally and are being used by the same application. **Improving scalability of the infrastructure creates heterogeneous environments that are more difficult to manage :** Even though organizations can achieve significant cost savings and increased flexibility of management by moving their business-critical applications into the cloud, this also creates a new environment that is fairly complex to monitor and manage. As a result, traditional IT management tools are not as effective in these environments as they are in managing the performance of internally hosted applications. This creates the challenge of finding a balance between scalability and flexibility of computing resources and ease of management and visibility into performance of the IT services relying on these resources.

Capabilities needed tools for measuring the impact of rules for assigning cloud resources on quality of end-user experience : One of the key benefits of cloud computing services is flexibility of assigning resources needed to support demand from business users. In order to achieve this benefit, many organizations deploying cloud computing services are defining rules for assigning cloud resources to each of their critical IT services and applications. However, the effectiveness of these policies depends on the visibility that organizations have into how cloud resources are being used. Organizations that have technology tools in place to monitor how changes in policies that control allocation of cloud computing resources impact the performance of business-critical applications, as measured from end-users' perspective, are more likely to reap the full benefits from the deployment of cloud computing.

An independent tool for monitoring/validating performance of a heterogeneous set of applications in the cloud : As organizations deploying cloud computing services trust third-party providers to deliver quality of service that would be acceptable to the end-users, they need to have technology tools in place to enable them to keep their service providers "honest" and have capabilities for monitoring levels of SLA achievements that go beyond monitoring capabilities provided by cloud vendors. As a part of their agreements with providers of public cloud services, organizations are requesting guarantees for levels of performance that service providers are expected to

deliver. However, in order to ensure that these service levels are met, organizations need to have independent monitoring tools in place that allow them to monitor not only actual levels of performance as experienced by business users.

Business Benefits

Organizations that are using the right mix of technology solutions for monitoring the performance of applications in the cloud are more likely to enjoy the following business benefits :

- Prevention and resolution of performance issues in a timely manner. Organizations that have visibility into resource utilization in the cloud are more likely to make educated and timely decisions about resource allocation and therefore, to prevent performance problems before they impact their business users.
- Ability to support changes in business demand. Full visibility into the performance of cloud services allows organizations to unlock the benefits of cloud computing, especially when it comes to improved flexibility of IT management. Organizations that have end-to-end visibility into the performance of cloud services and their internal infrastructure are able to make better decisions about adding or subtracting resources to support changes in business demand, which allows them to ensure a high level of quality of end-user experience at optimal cost.
- Ability to optimize spending decisions. Organizations deploying independent tools for monitoring performance, SLA achievements and usage of cloud services are more likely to be able to make educated decisions about the return they are getting from their investment in cloud services.

Q.21 What is the use of "EUCALYPTUS" in cloud computing?

Ans. Eucalyptus is a Linux-based open-source software architecture that implements efficiency-enhancing private and hybrid clouds within an enterprise's existing IT infrastructure.

Eucalyptus is an acronym for "Elastic Utility Computing Architecture for Linking Your Programs to Useful Systems."

A Eucalyptus private cloud is deployed across an enterprise's "on premise" data center infrastructure and is accessed by users over enterprise intranet. Thus, sensitive data remains entirely secure from external intrusion behind the enterprise firewall.

PREVIOUS YEARS QUESTIONS

PART-A

Q.1 *Why buffer is used in Amazon web services?*

Ans. In order to make system more efficient against the burst of traffic or load, buffer is used. It synchronises different component. The component always receives and processes the request in an unbalanced way. The balance between different components are managed by buffer and makes them work at same speed to provide faster services.

Q.2 *What is Amazon SQS?*

Ans. To communicate between different connectors Amazon SQS message is used, between various components of Amazon, it acts as a communicator.

Q.3 *Explain the security usage in Amazon web services model.*

Ans. The security usage in Amazon was services model :

- AWS supports security groups.
- Access is provided to create a security group for a jump box with SSH access only for port 22 open. Later a webserver group and a database groups are created.

Q.4 *Write any three key components of AWS.*

Ans. Three key components of AWS :

- **Identity and Access Management** : Improvised security and identity management are provided for an AWS account.
- **Simple E-mail Services** : Sending of e-mail is done by using a regular SMTP.
- **Route 53** : It is a DNS (Domain Name Server) web-based service platform.

Q.5 *Explain elastic compute cloud and elastic block stores.*

Ans. Elastic Compute Cloud : It allows on-demand computing resources for hosting applications and essential useful for unpredictable workloads.

Elastic Block Stores : They are storage volumes attached to elastic compute cloud and allow the data lifespan of a single elastic compute cloud.

Q.6 *What are the regions and availability zones in AWS?*

Ans. Regions : It is a geographical area which consist two or more availability zones. It is a collection of data centres which are completely isolated from other regions.

Availability Zones : It is a data centre that can be somewhere in the country or city.

PART-B

Q.7 *Compare Amazon, Azure and Google app Engine.*
[R.T.U. 2019]

Ans.

Vendor	Strengths	Weaknesses
AWS	<ul style="list-style-type: none"> • Dominant market position • Extensive, mature offerings • Support for large organizations • Extensive training • Global reach 	<ul style="list-style-type: none"> • Difficult to use • Cost management • Overwhelming options
Microsoft Azure	<ul style="list-style-type: none"> • Second largest provider • Integration with Microsoft tools and software • Broad feature set • Hybrid cloud • Support for open source 	<ul style="list-style-type: none"> • Issues with documentation • Incomplete management tooling
Google	<ul style="list-style-type: none"> • Designed for cloud-native businesses • Commitment to open source and portability • Deep discounts and flexible contracts • DevOps-expertise 	<ul style="list-style-type: none"> • Late entrant to IaaS market • Fewer features and services • Historically not as enterprise focused

Q.8 Explain cloud federation in details. [R.T.U. 2019]

Ans. Cloud Federation : Cloud federation is the practice of interconnecting the computing environments of two or more service providers for the purpose of load balancing traffic and accommodating spikes in demand.

Cloud federation requires one provider to wholesale or rent computing resources to another cloud provider. Those resources become a temporary or permanent extension of the buyer's cloud computing environment, depending on the specific federation agreement between providers.

Cloud federation offers two substantial benefits to cloud providers. First, it allows providers to earn revenue from computing resources that would otherwise be idle or underutilized. Second, cloud federation enables cloud providers to expand their geographic footprints and accommodate sudden spikes in demand without having to build new points-of-presence (POPs).

Service providers strive to make all aspects of cloud federation—from cloud provisioning to billing support systems (BSS) and customer support—transparent to customers. When federating cloud services with a partner, cloud providers will also establish extensions of their customer-facing service-level agreements (SLAs) into their partner provider's data centers.

Cloud Federation refers to the unionization of software, infrastructure and platform services from disparate networks that can be accessed by a client via the internet. The federation of cloud resources is facilitated through network gateways that connect public or external clouds, private or internal clouds (owned by a single entity) and/or community clouds (owned by several cooperating entities); creating a hybrid cloud computing environment. It is important to note that federated cloud computing services still rely on the existence of physical data centers.

The benefits of Cloud Federation are as follows :

The federation of cloud resources allows clients to optimize enterprise IT service delivery. The federation of cloud resources allows a client to choose the best cloud services provider, in terms of flexibility, cost and availability of services, to meet a particular business or technological need within their organization. Federation across different cloud resource pools allows applications to run in the most appropriate infrastructure environments. The federation of cloud resources also allows an enterprise to distribute workloads around the globe, move data between disparate networks and implement innovative security models for user access to cloud resources.

The federated cloud connects these local infrastructure providers to a global marketplace that enables each participant to buy and sell capacity on demand. As a provider, this gives you instant access to global infrastructure on an unprecedented scale. If your customer suddenly needs a few hundred new servers, you just buy the capacity they need from the marketplace. If a customer needs to accelerate a website or an application in Hong Kong, Tokyo or Latvia, you simply subscribe to those locations and make use of the infrastructure that's already there.

As part of a cloud federation, even a small service provider can offer a truly global service without spending a dime building new infrastructure. For companies with spare capacity in the data center, the federation also provides a simple way to monetize that capacity by submitting it to the marketplace for other providers to buy, creating an additional source of revenue.

Flexible : As your business changes or application evolves, you can easily reflect these changes in Amazon Simple DB without worrying about breaking a rigid schema or needing to refactor code - simply add another attribute to your Amazon Simple DB data set when needed. You can also choose between consistent or eventually consistent read requests, gaining the flexibility to match read performance and consistency requirements to the demands of your application, or even disparate parts within your application.

Simple to Use : Amazon Simple DB provides streamlined access to the store and query functions that traditionally are achieved using a relational database cluster - while leaving out other complex, often-unused database operations. The service allows you to quickly add data and easily retrieve or edit that data through a simple set of API calls.

Secure: Amazon Simple DB provides an https end point to ensure secure, encrypted communication between your application or client and your domain. In addition, through integration with AWS Identity and Access Management, you can establish user or group-level control over access to specific Simple DB domains and operations.

Inexpensive : Amazon Simple DB passes on to you the financial benefits of Amazon's scale. You pay only for resources you actually consume. For Amazon Simple DB, this means data store reads and writes are charged by complete resources consumed by each operation, and you aren't billed for compute resources when you aren't actively using them (i.e. making requests).

Q.11 Explain CRM.

[R.T.U. 2016]

Ans. CRM : In CRM (Customer Relationship Management), CRM software is a category of enterprise software that covers a broad set of applications and software designed to help businesses manage customer data and customer interaction, access business information, automate sales, marketing and customer support and also manage employee, vendor and partner relationships.

Customer relationship management (CRM) is a term that refers to practices, strategies and technologies that companies use to manage and analyze customer interactions and data throughout the customer lifecycle, with the goal of improving business relationships with customers, assisting in customer retention and driving sales growth. CRM systems are designed to compile information on customers across different channels or points of contact between the customer and the company which could

include the company's website, telephone, live chat, direct mail, marketing materials and social media. CRM systems can also give customer-facing staff detailed information on customers' personal information, purchase history, buying preferences and concerns.

Today's CRM Software

CRM software is designed to help businesses meet the overall goals of customer relationship management. Today's CRM software is highly scalable and customizable, allowing businesses to gain actionable customer insights with a back-end analytical engine, view business opportunities with predictive analytics, streamline operations and personalize customer service based on the customer's known history and prior interactions with your business.

CRM software is commonly used to manage a business-customer relationship, however CRM software systems are also used in the same way to manage business contacts, clients, contract wins and sales leads.

Q.12 Explain the data analysis application of cloud computing.

[R.T.U. 2018, 2015]

Ans. Data analysis of cloud computing works as following manners.

1. **Social Media:** A popular use for cloud data analytics is compounding and interpreting social media activity. Before cloud drives became practical, it was difficult processing activity across various social media sites, especially if the data was stored on different servers. Cloud drives allow for the simultaneous examination of social media site data so results can be quickly quantified and time and attention allocated accordingly.

2. **Tracking Products:** Long thought of as one of the kings of efficiency and forethought, it is no surprise Amazon.com uses data analytics on cloud drives to track products across their series warehouses and ship items anywhere as needed, regardless of items proximity to customers. Alongside Amazon's use of cloud drives and remote analysis, they are also a leader in big data analysis services thanks to their Redshift initiative. Redshift gives smaller organizations many of the same analysis tools and storage capabilities as Amazon and acts as an information warehouse, preventing smaller businesses from having to spend money on expensive hardware.

3. **Tracking Preference:** Over the last decade or so, Netflix has received a lot of attention for its DVD deliver service and the collection of movies hosted on their website. One of the highlights of their website is its movie recommendations, which tracks the movies users watch and recommends others they might enjoy, providing a

service to clients while supporting the use of their product. All user information is remotely stored on cloud drives so user's preferences do not change from computer to computer. Because Netflix retained all their user's preferences and tastes in movies and television, they were able to create a television show that statistically appealed to a large portion of their audience based on their demonstrated taste. Thus in 2013, Netflix's House of Cards became the most successful internet television series ever, all thanks to their data analysis and information stored on clouds.

4. **Keeping Records:** Cloud analytics allows for the simultaneous recording and processing of data regardless of proximity to local servers. Companies can track the sales of an item from all their branches or franchises across the United States and adjust their production and shipments as necessary. If a product does not sell well, they do not need to wait for inventory reports from area stores and can instead remotely manage inventories from data automatically uploaded to cloud drives. The data stored to clouds helps to make business run more efficiently and gives companies a better understanding of their customers behavior.

Q.13 What do you mean by SLA? Explain elements of SLA.

[R.T.U. 2017]

Ans. Key Elements of SLA

1. **Context Setting Information :** Context Setting Information "sets the table" for the SLA. It explains the purpose and scope of the agreement, the parties involved, and the agreement's underlying assumptions.

2. **Description of Services :** The Description of Services also helps "set the table". It discusses the services provided and the services not provided. This helps clarify things in cases where agents might assume the availability of certain services. This section needs to be clear and specific.

3. **Service Standards :** Including Service Standards ensures that both parties share a common understanding about the conditions under which the stated services will be provided.

4. **Service Tracking/Reporting :** Service Tracking/Reporting information identifies how service effectiveness will be assessed and communicated. Tracking/reporting information includes a good set of metrics. They should reflect the project's major objectives.

5. **Periodic Review :** This section ensures ongoing communication between the two parties and formal systematic attention to service adequacy. When things are going well managers sometimes fail to hold periodic reviews. Hold reviews at the appropriate times even when things are going well.

6. **Change Process of Service Level Agreement :** The Change Process section provides a formal tool for modifying the agreement to address changing service needs and priorities. Accordingly, a good SLA provides a mechanism for periodic reviews and modifications as needed.

A Service Level Agreement can include more than these 6 elements. But all effective SLAs include at least these 6 elements.

Q.14 Write short note on account hijacking.

[R.T.U. 2017]

Ans. Account Hijacking : Cloud account hijacking occurs when an individual or organization's cloud account is stolen or hijacked by an attacker. Cloud account hijacking is a common tactic in identity theft schemes. The attacker uses the stolen account information to conduct malicious or unauthorized activity. When cloud account hijacking occurs, an attacker typically uses a compromised email account or other credentials to impersonate the account owner.

In a report from the Cloud Security Alliance service traffic hijacking was identified as the third-greatest cloud computing security risk. These types of security breach occur when attackers hijack cloud accounts by stealing security credentials and eavesdropping on activities and transactions. Attackers manipulate data, insert false information, and redirect clients to illegitimate sites.

Cloud account hijacking at the enterprise level can be particularly devastating, depending on what the attackers do with the information. Company integrity and reputations can be destroyed, and confidential data can be leaked or falsified causing significant cost to businesses or their customers. Legal implications are also possible for companies and organisations in highly regulated industries, such as healthcare, if clients' or patients' confidential data is exposed during cloud account hijacking incidents.

Q.15 Write short note on cloud application platform - Integration of private and public cloud.

[R.T.U. 2016]

Q.17 Explain the working of MAP Reduce.

/R.T.U. 2019/

OR
Explain Google APP engine. (R.T.U. 2018)

What is Google App. Explain the architecture of Google App Engine in detail.
[R.T.U. 2015]

Explain Google App Engine in detail with architecture.

Ans: Google App Engine is a Platform as a Service (PaaS) offering that lets you build and run applications on Google's infrastructure. App Engine applications are easy to build, easy to maintain and easy to scale as our traffic and data storage needs change. With App Engine, there are no servers for us to maintain. We simply upload our application and it's ready to go.

App Engine Runtime Environment

Google App Engine supports apps written in a variety of programming languages.

1. **Java:** Using App Engine's Java runtime environment, we can build our application using standard Java technologies.
 2. **Python:** App Engine features a fast Python interpreter and standard Python libraries.
 3. **PHP:** App Engine uses Google's Cloud Platform services under the hood when we call standard PHP functions.
 4. **Go:** App Engine features a Go runtime environment that runs natively compiled Go code.
- Google App Engine makes it easy to build and deploy an application that runs reliably even under heavy load and with large amounts of data. It includes the following features:
1. Persistent storage with queries, sorting and transactions.
 2. Automatic scaling and load balancing.
 3. Asynchronous task queues for performing work outside the scope of a request.
 4. Scheduled tasks for triggering events at specified times or regular intervals.

5. **Integration with other Google cloud services and APIs.** Applications run in a secure, sandboxed environment, allowing App Engine to distribute requests across multiple servers and scaling servers to meet traffic demands. Your application runs within its own secure, reliable environment that is independent of the hardware, operating system or physical location of the server.

1. All of the APIs and libraries available to App Engine.
2. A simulated, secure sandbox environment, that emulates all of the App Engine services on your local computer.
3. Deployment tools that allow you to upload your application to the cloud and manage different versions of your application.

The SDK manages your application locally, while the Google Developers Console manages your application in production. The Developers Console uses a web-based interface to create new applications, configure domain names, change which version of our application is live, examine access and error logs and much more.

Limits : App Engine gives you 1 GB of data storage and traffic for free, which can be increased by enabling paid applications. However, some features impose limits unrelated to quotas to protect the stability of the system.




Fig. : Google App Engine Architecture

Q.18 Describe some example of CRM and ERP based on cloud computing. [R.T.U. 2019]

Ans. Examples of CRM solutions : CRM software typically falls into four broad categories: outsourced, off-the-shelf, bespoke and managed solutions.

Outsourced CRM solutions : This includes web-based CRM solutions for your business, including cloud CRM software. A few examples of outsourced CRM are:

- Salesforce
- SugarCRM
- Really Simple Systems
- NetSuite CRM
- Zoho

This approach is ideal if you need to implement a solution quickly and you don't have the in-house skills to tackle the job from scratch. It is also a good solution if you are already geared towards e-commerce.

Off-the-shelf CRM solutions : Several software companies offer CRM applications that integrate with existing packages. Cut-down versions of such software may be suitable for smaller businesses. One such example is Microsoft Dynamics CRM.

Off-the-shelf products are generally the cheapest option as you are investing in standard software components. The downside is that the software may not always do precisely what you want and you may have to trade off functionality for convenience and price. The key is to be flexible without compromising too much.

Bespoke CRM software : Consultants and software engineers can customise or create a CRM system and integrate it with your existing software. However, this can be expensive and time-consuming. If you choose the custom-tailored option, make sure that you carefully specify exactly what you want. Costs will vary, so it may be worth getting several quotes from different reputable professionals.

Managed CRM solutions : A half-way house between bespoke and outsourced solutions, this involves renting a customised suite of CRM applications as a bespoke package. This can be cost effective but it may mean that you have to compromise in terms of functionality. Examples include **Maximizer CRM** and **Sage**.

ERP System Examples

Episcor ERP 10: Episcor's flexible collection of ERP tools can be deployed on company computers, in the cloud or on an external server. This system highlights collaboration by allowing users to access data from different devices.

companies a 103 percent return on investment within about a year and a half.

IFS Full Suite ERP : If you're looking for a user-friendly interface for project management, the IFS Full Suite ERP solution is ideal. This system is perfect for manufacturing and engineering companies

Infor ERP Syteline : Based in the cloud, this ERP solution can be customized for large-scale manufacturers and distributors. It's scalable, so it easily adapts as your business grows.

Microsoft Dynamics AX 2012 : International businesses will find Microsoft's ERP solution useful for most of their departments. The software is powerful and easy to implement, improving its value for driving organization and productivity.

Oracle JD Edwards EnterpriseOne : Oracle's most recent product is ideal for corporations seeking a financially focused ERP solution. This ERP automates a majority of a company's accounting, whether for small projects or overall evaluation. EnterpriseOne helps businesses amplify their financial success.

Sage ERP X3 : For small-to-mid sized companies, Sage ERP X3 improves productivity within fundamental business systems. The software has complete mobile capability, making it ideal for use in a flexible corporate environment.

SAP Business ByDesign : Another cloud-based ERP tool, SAP Business ByDesign can manage any facet of a corporation with around-the-clock availability.

SYSPRO 7 : Completely customizable, **SYSPRO 7** can be incorporated into any company's system management. Whether your needs are based on financials, customer service, sales, distribution or manufacturing, you can personalize **SYSPRO 7** for your company.

Q.19 Describe the amazon EC2 and its features.

[R.T.U. 2019]

OR

Explain the architecture of Amazon EC2.

AWS Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers.

Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction.

It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate themselves from common failure scenarios.

To use the EC2, a subscriber creates an Amazon Machine Image (AMI) containing the operating system, application programs and configuration settings. Then the AMI is uploaded to the Amazon Simple Storage Service (Amazon S3) and registered with Amazon EC2, creating a so-called AMI identifier (AMI ID). Once this has been done, the subscriber can requisition virtual machines on an as-needed basis. Capacity can be increased or decreased in real time from as few as one to more than 1000 virtual machines simultaneously. Billing takes place according to the computing and network resources consumed.

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

Features of Amazon EC2

Amazon EC2 provides the following features:

- Virtual computing environments, known as *instances*
- Preconfigured templates for your instances, known as *Amazon Machine Images (AMIs)*, that package the bits you need for your server (including the operating system and additional software)
- Various configurations of CPU, memory, storage and networking capacity for your instances, known as *instance types*
- Secure login information for your instances using *key pairs* (AWS stores the public key and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as *instance store volumes*

- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as *Amazon EBS volumes*
- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as *regions* and *Availability Zones*
- A firewall that enables you to specify the protocols, ports and source IP ranges that can reach your instances using *security groups*
- Static IP addresses for dynamic cloud computing, known as *Elastic IP addresses*
- Metadata, known as *tags*, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud and that you can optionally connect to your own network, known as *virtual private clouds (VPCs)*.

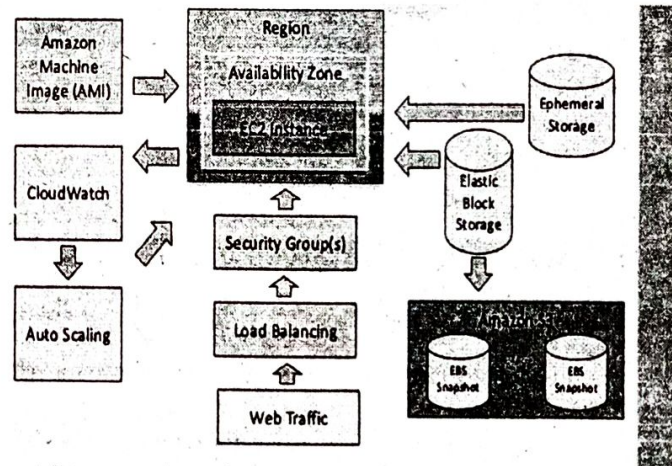


Fig. : Amazon EC2 Architecture

Amazon Elastic Compute Cloud (EC2) it is a hardware as a Service. A Web service that provides resizable compute capacity in the cloud. It is designed to make Web-scale computing easier for developers. It is in reality Xen virtual machine instances running on AMD x86; each instance has 2GB RAM and 150GB disk. A simple Web service interface that provides complete control of your computing resources.

Benefits of EC2:

- Reduces the time required to obtain and boot new server instances to minutes
- Quickly scales capacity, both up and down, as your computing requirements change
- Changes the economics of computing:
 - (i) Pay only for capacity that you actually use
 - (ii) $a + bc$ becomes just bc

1. **Execution and Operation** : This phase includes the actual start of setting up the cloud services, populating the respective cloud service with relevant data, on boarding and training users, setting up communication channels and further operational activities while using the respective cloud services.
2. **Updates and Amendments** : This phase includes updated or otherwise amended needs, goals and assumptions by the Cloud Service Customer during the term of the on-going cloud services arrangements, as well as improved or added cloud services by the CSP there under. It also includes optimisation of the respective cloud services by CSP as per (contractual or other) non-compliance, breaches and other incidents during that term.
3. **Escalation** : This phase deals with contractual or other and non-compliance, breaches and other incidents during the term of the ongoing cloud services arrangements that have resulted in a dispute that needs escalation, (perhaps even litigation as a last resort), negotiation and resolution.

Termination of Cloud Service Lifecycle : Why is it important? You should already think about termination in phase 1, as an SLA can be used to arrange the conditions under which the Cloud customer's data (including but not limited to for instance Personal Identifiable Information or PII) will be exported and returned to the cloud customer, and not retained by the cloud service provider (to the extent mandatorily possible).

Termination and Consequences of Termination

This phase deals with the end of the relationship between CSP and CSC, including the end of the legal relationship even though the latter will generally continue for several years after any termination as per mandatory laws and legislation. This last phase for instance includes the assessment of alternatives, settlement and termination arrangements, cloud services transition projects and services, data export, customer and (end)use care and diligence, and adequate data deletion.

Q.21 Write short notes on Aneka.

Ans. Aneka : Aneka is a market oriented Cloud development and management platform with rapid application development and workload distribution capabilities. Aneka is an integrated middleware package which allows you to seamlessly build and manage an interconnected network in addition to accelerate

development, deployment and management of distributed applications using Microsoft .NET frameworks on these networks. It is market oriented since it allows you to build, schedule, provision and monitor results using pricing, accounting, QoS/SLA services in private and/or public (leased) network environments.

Aneka is a workload distribution and management platform that accelerates applications in Microsoft .NET framework environments. Some of the key advantages of Aneka over other GRID or Cluster based workload distribution solutions include:

Aneka is a platform and a framework for developing distributed applications on the Cloud. It harnesses the spare CPU cycles of a heterogeneous network of desktop PCs and servers or datacenters on demand. Aneka provides developers with a rich set of APIs for transparently exploiting such resources and expressing the business logic of applications by using the preferred programming abstractions. System administrators can leverage on a collection of tools to monitor and control the deployed infrastructure. This can be a public cloud available to anyone through the Internet or a private cloud constituted by a set of nodes with restricted access.

The Aneka based computing cloud is a collection of physical and virtualized resources connected through a network, which are either the Internet or a private intranet. Each of these resources hosts an instance of the Aneka Container representing the runtime environment where the distributed applications are executed. The container provides the basic management features of the single node and leverages all the other operations on the services that it is hosting. The services are broken up into fabric, foundation and execution services. Fabric services directly interact with the node through the Platform Abstraction Layer (PAL) and perform hardware profiling and dynamic resource provisioning. Foundation services identify the core system of the Aneka middleware, providing a set of basic features to enable Aneka containers to perform specialized and specific sets of tasks. Execution services directly deal with the scheduling and execution of applications in the Cloud.

One of the key features of Aneka is the ability of providing different ways for expressing distributed applications by offering different programming models; execution services are mostly concerned with providing the middleware with an implementation for these models. Additional services such as persistence and security are transversal to the entire stack of services that are hosted by the container. At the application level, a set of different components and tools are provided to: (1) simplify the development of applications (SDK); (2) porting existing

applications to the Cloud; and (3) monitoring and managing the Aneka Cloud.

A common deployment of Aneka is presented at the side. An Aneka based Cloud is constituted by a set of interconnected resources that are dynamically modified according to the user needs by using resource virtualization or by harnessing the spare CPU cycles of desktop machines. If the deployment identifies a private Cloud all the resources are in house, for example within the enterprise. This deployment is extended by adding publicly available resources on demand or by interacting with other Aneka public clouds providing computing resources connected over the Internet.

- Rapid deployment tools and framework,
- Ability to harness multiple virtual and/or physical machines for accelerating application result
- Provisioning based on QoS/SLA
- Support of multiple programming and application environments
- Simultaneous support of multiple run-time environments
- Built on-top of Microsoft .NET framework, with support for Linux environments through Mono

Q.22 Define Data Shredding.

Ans. Data Shredding : Shredding is a process of irreversible file destruction, so that its contents could not be recovered. Sometimes the same process is referred as erasing or wiping; we prefer to call it shredding in an analogy with paper shredding machines, which are used for disposing sensitive documents.

Importance of Shredding

"What do I need shredding for? I want to encrypt my file, not to destroy it." This is a rather common question and the answer will be obvious if we look at a simple analogy.

Imagine that you need to encrypt a paper-written message using the traditional James Bond way. The process is fairly straightforward – encrypt the message word by word, writing the result down on a separate paper sheet. When the whole message is encrypted, burn the sheet with the original message.

We can erase each word after it is encrypted and write the result onto the erased space. When all the words are encrypted, you will have the same paper sheet on which the original message has been replaced with the encrypted text.

One does not need to be an expert in order to see disadvantages of this method. It takes more time, requires a lot of extra work and what is worst, the original message

Cloud Computing

1. Protection against Hardware Failures :

Because every application is made up of multiple instances of each role, hardware failures – a disk crash, a network fault or the death of a server machine won't take down the application. To help with this, the fabric controller doesn't choose machines for an application's instances at random. Instead, different instances of the same role are placed in different fault domains. A fault domain is a set of hardware computers, switches and more, that share a single point of failure. (For example, all of the computers in a single fault domain might rely on the same switch to connect to the network.) Because of this, a single hardware failure can't take down an entire application. The application might temporarily lose some instances, but it will continue to behave correctly.

2. Protection Against Software Failures : Along with hardware failures, the fabric controller can also detect failures caused by software. If the code in an instance crashes or the VM in which it's running goes down, the fabric controller will start either just the code or, if necessary, a new VM for that role. While any work the instance was doing when it failed will be lost, the new instance will become part of the application as soon as it starts running.

3. The ability to Update Applications with No Application Downtime : Whether for routine maintenance or to install a whole new version, every application needs to be updated. An application built using the Windows Azure programming model can be updated while it's running there's no need to take it down.

Q.24 Discuss about Amazon AWS services.

Ans. Amazon Web Services (AWS) is a comprehensive, evolving cloud computing platform provided by Amazon.com. Web services are sometimes called cloud services or remote computing services. The first AWS offerings were launched in 2006 to provide online services for websites and client-side applications.

To minimize the impact of outages and ensure robustness of the system, AWS is geographically diversified into regions. These regions have central hubs in the Eastern USA, Western USA (two locations), Brazil, Ireland, Singapore, Japan and Australia. Each region comprises multiple smaller geographic areas called availability zones.

Amazon Web Services or AWS, is a cloud computing platform from Amazon that provides customers with a wide array of cloud services. Among the cloud options offered by Amazon AWS are Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon Virtual Private Cloud

(Amazon VPC), Amazon SimpleDB and Amazon WorkSpaces.

Amazon first debuted its Amazon Web Services in 2006 as a way to enable the use of online services by client-side applications or other web sites via HTTP, REST or SOAP protocols. Amazon bills customers for Amazon AWS based on their usage of the various Amazon Web Services.

In 2012, Amazon launched the AWS Marketplace to accommodate and grow the emerging ecosystem of AWS offerings from third-party providers that have built their own solutions on top of the Amazon Web Services platform. The AWS Marketplace is an online store for Amazon Web Services customers to find, compare and begin using AWS software and technical services.

The Differences that Distinguish AWS from Others:

AWS is readily distinguished from other vendors in the traditional IT computing landscape because it is:

- 1. Flexible :** AWS enables organizations to use the programming models, operating systems, databases and architectures with which they are already familiar. In addition, this flexibility helps organizations mix and match architectures in order to serve their diverse business needs.
- 2. Cost-effective :** With AWS, organizations pay only for what they use, without up-front or long-term commitments.
- 3. Scalable and Elastic :** Organizations can quickly add and subtract AWS resources to their applications in order to meet customer demand and manage costs.
- 4. Secure :** In order to provide end-to-end security and end-to-end privacy, AWS builds services in accordance with security best practices, provides the appropriate security features in those services and documents how to use those features.
- 5. Experienced :** When using AWS, organizations can leverage Amazon's more than fifteen years of experience delivering large-scale, global infrastructure in a reliable, secure fashion.

The growing AWS collection offers over three dozen diverse services including:

- CloudDrive, which allows users to upload and access music, videos, documents and photos from Web-connected devices. The service also enables users to stream music to their devices.
- CloudSearch, a scalable search service typically used to integrate customized search capabilities into other applications.
- Dynamo Database (also known as DynamoDB or DDB), a fully-managed NoSQL database service known for low latencies and scalability.

CLC.89

CLC.90

- Elastic Compute Cloud, which allows business subscribers to run application programs and can serve as a practically unlimited set of virtual machines (VMs).
 - ElastiCache, a fully managed caching service that is protocol-compliant with Memcached, an open source, high-performance, distributed memory object caching system for speeding up dynamic Web applications by alleviating database load.
 - Mechanical Turk, an application program interface (API) that allows developers to integrate human intelligence into Remote Procedure Calls (RPCs) using a network of humans to perform tasks that computers are ill-suited for.
 - RedShift, a petabyte-scale data warehouse service designed for analytic workloads, connecting to standard SQL-based clients and business intelligence tools.
 - Simple Storage Service (S3), a scalable, high-speed, low-cost service designed for online backup and archiving of data and application programs.
- All AWS offerings are billed according to usage.

Q.25 Explain most common applications of cloud computing.

Ans. Cloud computing applications or apps, are the cloud-based services also known as Software as a Service (SaaS). Programs that once had to be installed on computers individually are now offered online and the only thing a person needs to access the program is an account and password. These apps can do everything from keeping track of notes to accounting. For both software providers and users there are a number of benefits to using cloud based apps:

A. Collaboration: Cloud apps give employees access to their information from anywhere around the globe. All you need is an Internet connection. This allows more collaborative working as multiple people can view and edit the same information at once, ensuring your team works efficiently.

B. Automatic Updates: Software as a service (SaaS) allows companies to ensure all users of their application are on the same version of the software. This is because they can provide automatic updates to cloud applications, rather than waiting for users to do it themselves. This also helps with support, as the company will know what version of the software is being used when issues are logged.