UNIT III

3 Virtualization Infrastructure and Docker

Syllabus

Desktop Virtualization - Network Virtualization - Storage Virtualization - System-level of Operating Virtualization - Application Virtualization - Virtual clusters and Resource Management - Containers vs. Virtual Machines - Introduction to Docker - Docker Components - Docker Container - Docker Images and Repositories.

Contents

3.1 Desktop Virtualization
3.2 Network Virtualization
3.3 Storage Virtualization
3.4 System - Level of Operating Virtualization
3.5 Application Virtualization
3.6 Virtual Clusters and Resource Management
3.7 Introduction to Docker
3.8 Two Marks Questions with Answers

3.1 Desktop Virtualization

- Desktop virtualization is a technology that allows the creation and storage of multiple user desktop instances on a single host, residing in a data center or the cloud. It is achieved by using a hypervisor, which resides on top of the host server hardware to manage and allow virtual desktops to utilize the computing power of the underlying server hardware.
- Fig. 3.1.1 shows desktop virtualization.

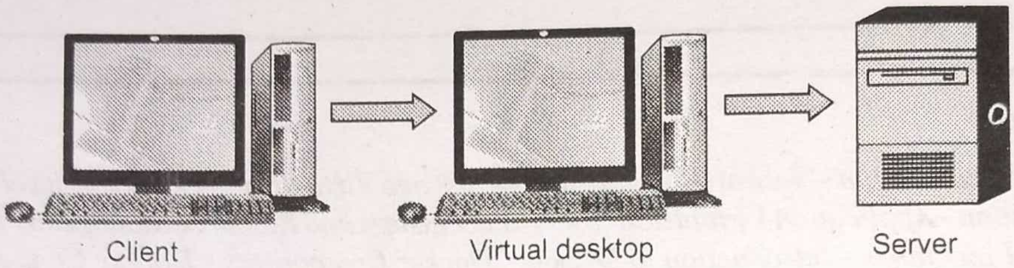

Fig. 3.1.1 Desktop virtualization

- The three most popular types of desktop virtualization are Virtual Desktop Infrastructure (VDI), Remote Desktop Services (RDS) and Desktop-as-a-Service (DaaS).

3.1.1 Types of Desktop Virtualization

1. Virtual desktop infrastructure

- A popular type of desktop virtualization is Virtual Desktop Infrastructure (VDI). VDI uses a VM to deliver persistent or non-persistent virtual desktops to many types of connected devices.
- With a persistent virtual desktop, each user has a unique, dedicated desktop image they can customize with apps and data, knowing the desktop will be saved for future use.
- A non-persistent VDI allows users to access a virtual desktop from an identical pool when they need it. Once the user logs out of a non-persistent VDI, the VDI reverts to its unaltered state.
- Characteristics of VDI :
  i) Virtual desktops live within virtual machines on a centralized server.
  ii) Each virtual desktop includes an operating system image, typically Microsoft Windows.
  iii) The virtual machines are host-based, meaning multiple instances of them can be housed on the same server within the datacenter.

iv) End clients, such as PCs, tablets or thin client terminals, must be constantly connected to the centrally managed server so they can maintain access to the virtualized desktops they're hosting.

v) The connection broker is a software layer that acts as an intermediary between users and virtual resources, which finds a virtual desktop within the resource pool for each client upon successful access of the VDI environment.

- Here are some reasons why VDI is beneficial :
  a) Save money on licensing and individual Workstations/PCs by using thin clients.
  b) Fully secured virtual environment that is fully monitored and managed.
  c) Centralized management and backups.
  d) Secure remote access from anywhere in the world.
  e) Cost reduction for multiple software licenses.
- **Disadvantages :**
  a) If an individual requires different applications from the other users, they will require a completely different image, without changing the applications for other users.
  b) A substantial initial outlay is required for the main server hardware, storage and network infrastructure. This might not be feasible for some smaller businesses.
  c) Administrators, savvy to the limitations, problem solving and installation of VDIs will either have to be brought in or existing IT staff given the relevant training.
  d) If a problem occurs, this will generally affect all users, rather than being able to isolate problems if operating systems run off individual PCs.

2. Remote Desktop Services

- Remote Desktop Services (RDS) or Remote Desktop Session Host (RDSH) are beneficial where only limited applications require virtualization. They allow users to remotely access Windows applications and desktops using the Microsoft Windows server operating system.
- RDS is a more cost-effective solution, since one Windows server can support multiple users.

3. Desktop-as-a-Service (DaaS)

- Desktop-as-a-service (DaaS) is a flexible desktop virtualization solution that uses cloud-based virtual machines backed by a third-party provider. Using DaaS, organizations can outsource desktop virtualization solutions that help a user to access computer applications and desktops from any endpoint platform or device.

3.1.2 Benefits of Desktop Virtualization

1. Resource utilization : Since IT resources for desktop virtualization are concentrated in a data center, resources are pooled for efficiency.
2. Remote workforce enablement : Since each virtual desktop resides in central servers, new user desktops can be provisioned in minutes and become instantly available for new users to access.
3. VDI offers security improvements compared with running everything locally.

3.2 Network Virtualization

- Network virtualization refers to the technology that enables partitioning or aggregating a collection of network resources and presenting them to various users in a way that each user experiences an isolated and unique view of the physical network.
- Network virtualization creates virtual networks whereby each application sees its own logical network independent of the physical network.
- A virtual LAN (VLAN) is an example of network virtualization that provides an easy, flexible, and less expensive way to manage networks.
- VLANs make large networks more manageable by enabling a centralized configuration of devices located in physically diverse locations.
- Fig. 3.2.1 shows network virtualization.
- Consider a company in which the users of a department are separated over a metropolitan area with their resources centrally located at one office.
- In a typical network, each location has its own network

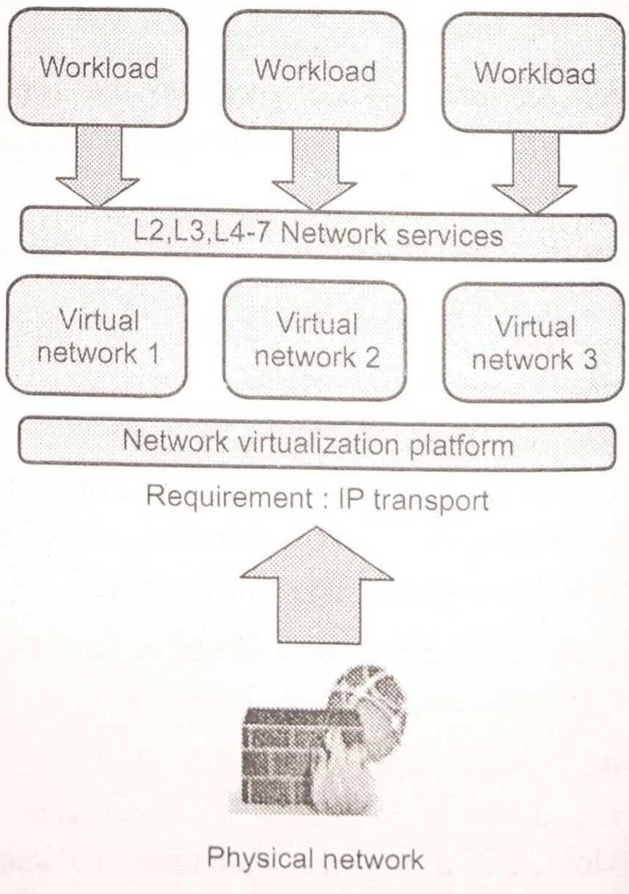

Fig. 3.2.1 Network virtualization

connected to the others through routers. When network packets cross routers, latency influences network performance.

- With VLANs, users with similar access requirements can be grouped together into the same virtual network. This setup eliminates the need for network routing.
- As a result, although users are physically located at disparate locations, they appear to be at the same location accessing resources locally.
- In addition to improving network performance, VLANs also provide enhanced security by isolating sensitive data from the other networks and by restricting access to the resources located within the networks.
- Network virtualization decouples the roles of the traditional Internet Service Providers (ISPs) into Infrastructure Providers (InPs) and Service Providers (SPs)
- Benefits :
  1. Reduces the number of physical devices needed.
  2. Easily segment networks.
  3. Permits rapid change / scalability and agile deployment.
  4. Security from destruction of physical devices.

3.3 Storage Virtualization

- Storage virtualization is a major component for storage servers, in the form of functional RAID levels and controllers. Operating systems and applications with device can access the disks directly by themselves for writing.
- Storage virtualization in cloud computing pools multiple physical storage arrays from Storage Area Networks (SANs) and makes them appear as a single virtual storage device. Virtualization storage separates the storage management software from the underlying hardware infrastructure to provide more flexibility and scalable pools of storage resources.
- Fig. 3.3.1 shows storage virtualization.
- Storage virtualization refers to the abstraction of storage systems from applications or computers. It is a foundation for the implementation of other technologies, such as thin provisioning and data protection, which are transparent to the server.
- Storage virtualization provides the ability to pool storage systems into a consolidated, shared capacity that can be managed from a central point of control.
- Example of storage virtualizations are host-based volume management, LUN creation, tape storage virtualization and disk addressing.
- Storage virtualization has the following characteristics :
  1. The availability of logical volumes separate from physical hard disk constraints.

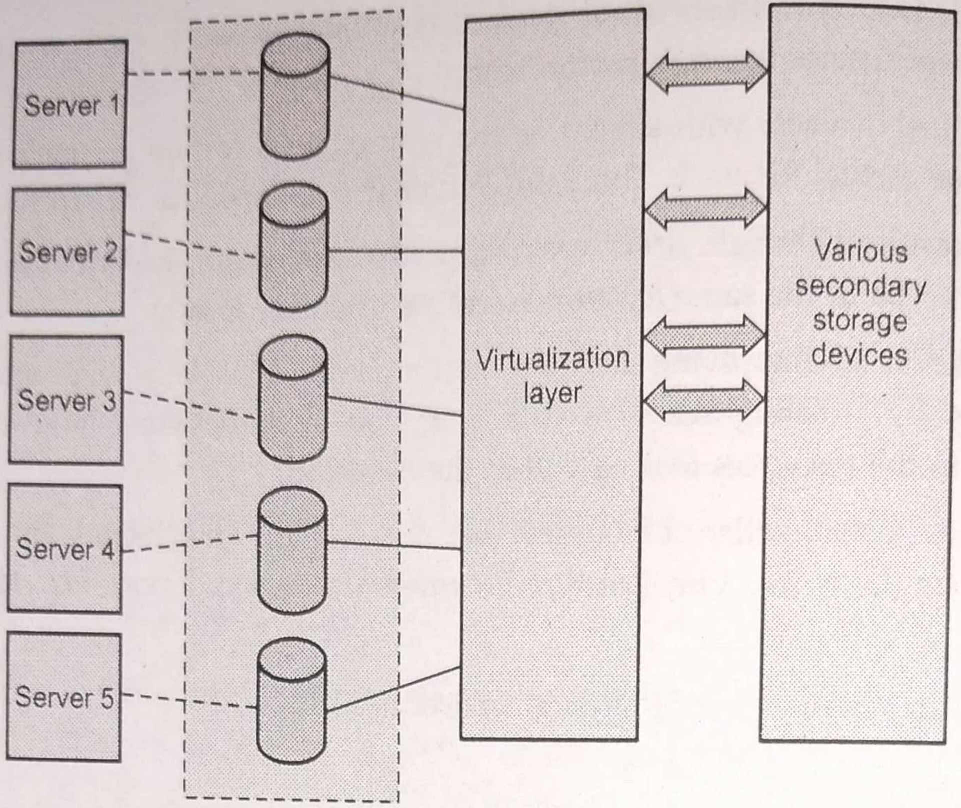

Fig. 3.3.1 Storage virtualization

2. The capability of abstracting multivendor storage devices into one group and reallocating storage space independently of size or physical location.
3. The capability of having automated storage optimization and management.

- Top level servers assigned one virtual volume, which is currently in use by an application. These virtual volumes are mapped to the actual storage in the arrays. When an I/O is sent to a virtual volume, it is redirected through the virtualization at the storage network layer to the mapped physical array.
- Primary types of storage virtualizations are block level virtualization and file virtualization.
- Currently there are three methods of storage virtualization :
  1. Server-based virtualization : This method places a management program on the host system and has the benefit of leveraging the SAN asset as it is.
  2. Fabric-based virtualization : This can be done via network switches or appliance servers. In both instances, independent appliances, such as switches, routers and dedicated servers are placed between servers and storage and have a storage virtualization function. The purpose behind this is to reduce the impact on the existing SAN and servers.
  3. Storage array-based virtualization : This is a virtualization implemented at the storage-system level.

3.3.1 Storage Virtualization Challenges

- Storage virtualization has evolved at a time when data explosion threatened to throw enterprise storage management totally out of gear.
- Traditionally, managing disk storage was once simple : If enterprises needed more space, they got a bigger disk drive. However, as data storage needs grew, multiple disk drives had to be added. Over time technologies such as RAID, network-attached storage and storage-area networks evolved to tackle these storage challenges.
- But managing and maintaining thousands of disk drives presented an even more serious challenge and storage virtualization emerged to tackle these.
  1. Scalability : Ensure storage devices perform appropriate requirements. Each array is managed independently.
  2. Functionality : Virtualized environment must provide same or better functionality. It must be continue to leverage existing functionality on arrays.
  3. Manageability : Virtualization device breaks end-to-end view of storage infrastructure and must integrate existing management tools.
  4. Support : Interoperability in multi-vendor environment.
- A good storage virtualization solution should :
  1. Enhance the storage resources it is virtualizing through the aggregation of services to increase the return of existing assets.
  2. Not add another level of complexity in configuration and management.
  3. Improve performance rather than act as a bottleneck in order for it to be scalable. Scalability is the capability of a system to maintain performance linearly as new resources are added.
  4. Provide secure multi-tenancy so that users and data can share virtual resources without exposure to other users' bad behavior or mistakes.
  5. Not be proprietary, but virtualize other vendor storage in the same way as its own storage to make the management seamless.

3.3.2 Types of Storage Virtualization

- Storage virtualization provides the ability to pool storage systems into a consolidated, shared capacity that can be managed from a central point of control. Virtualization can be implemented in both storage area network and network attached storage.
- Storage virtualization are of two types : Block level and File level.

3.3.3 Block Level Virtualization

- Block level virtualization is used in storage area network. The act of applying virtualization to one or more block-based storage services for the purpose of providing a new block service to clients. Some examples of block virtualization are disk aggregation.
- Block which is used for data storage is progression of bytes and bits and is made up of a proposed length. Data which is aligned in these blocks is called as blocked and inserting data into the block is called blocking.
- Block level storage virtualization provides storage to operating systems and applications in the form of virtual disks. Fig. 3.3.2 Shows block level virtualization.

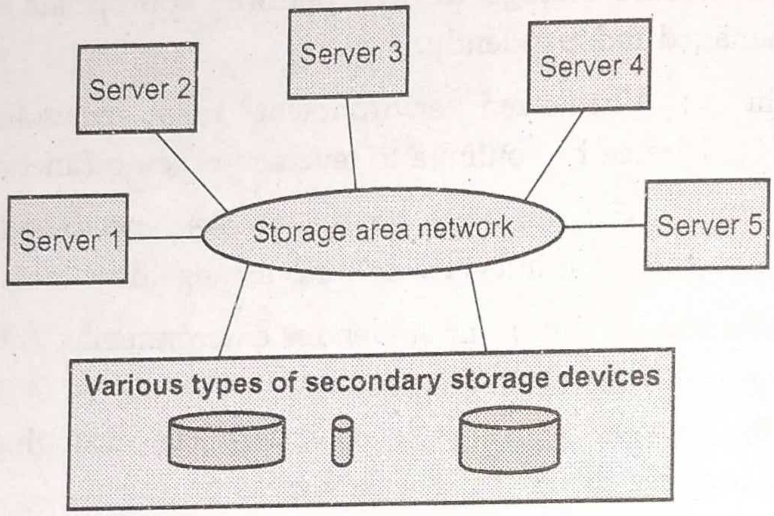

Fig. 3.3.2 Block level virtualization

- There are two types of block level virtualization. One is disk level virtualization, whereby an abstraction process moves data from a physical disk level to a LUN level and is presented as though it were a physical device.
- Another method is storage level virtualization, which, unlike disk level virtualization, hides the physical layer of RAID controllers and disks and hides and virtualizes the entire storage system.
- SCSI commands are transmitted in between the initiator and target. There is no overhead file system like an ext3.
- Block level file system utilizes FC, iSCSI and FCOE protocol.
- Block level file storage is pretty expensive but is very much reliable. It is highly customizable storage and is versatile and speedy.
- Block-level virtualization is usually just called storage virtualization and serves applications such as database software that need block-level access to data. The disks will typically (but not always) reside in Storage Area Network arrays (SANs).

3.3.4 File Level Virtualization

- Network attached storage uses file level virtualization.
- File level storage virtualization provides storage volumes to operating systems and applications in the form of files and directories. Access to storage is through network protocols, such as Common Internet File System and Network File Systems.
- Storage resources and capacity is may be underutilized because files are bound to a specific file server. It is necessary move the file from one server to another server.
- File-level storage is the predominant storage technology used on hard drives, Network-Attached Storage (NAS) systems and similar storage systems. Fig. 3.3.3 shows file level virtualization.

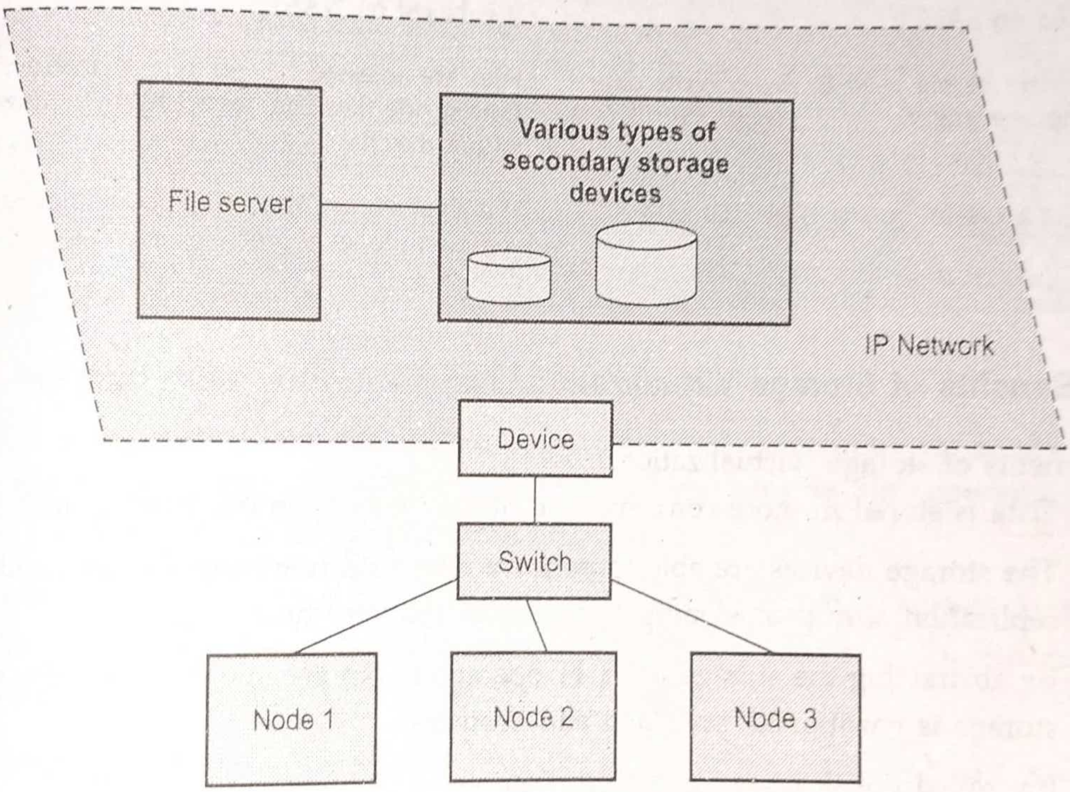

Fig. 3.3.3 File level virtualization

- Moving large number of files in not possible because it requires the server is to be down. Server and some applications need to be reconfigured with the new path. It creates the problem for network administrators for improving the storage efficiency while maintaining the required service level.
- This file of virtualization only simplifies the file mobility. It provides location transparency to user. File level storage works with an ext3 file system. Data is written and read into files, which have variable lengths.

- File level storage will not support of virtual machine file system. It support external boot up, which is essential for ESX and ESXi host servers.
- This type storage cannot handle heavy traffic on the network. Recovery of files is much faster in this level of data storage system. Storage resources and capacity is may be underutilized because files are bound to a specific file server. It is necessary move the file from one server to another server.

3.3.5 Difference between Block Level and File Level Virtualization

Block level	File level
Block-Level virtualization works before the file system exists. It replaces controllers and takes over at the disk level.	The server that uses the storage must have software installed on it in order to enable file-level usage.
It is based on SAN.	It is based on NAS.
Block addresses are used to Read/Write data to the storage media.	Files are accessed by "semantics" instructions. Data inside files is accessed by byte-ranges within the file.
Storage is accessible using fibre channel or iSCSI.	File level storage is usually accessible using common file level protocols such as CIFS and NFS.

3.3.6 Benefits of Storage Virtualization

- Benefits of storage virtualization :
  1. Data is stored in more convenient locations away from the specific host.
  2. The storage devices are able to perform advanced functions like de-duplication, replication, thin provisioning and disaster recovery functionality.
  3. By abstracting the storage level, IT operations can become more flexible in how storage is partitioned, provided and protected.
  4. Improved physical resource utilization.
  5. Lower total cost of ownership : Virtualized storage allows more to be done with the same or less storage.

3.4 System - Level of Operating Virtualization

- Operating - system - level virtualization is a server-virtualization method where the kernel of an operating system allows for multiple isolated user-space instances, instead of just one. Such instances, which are sometimes called containers and software containers.
- This refers to an abstraction layer between traditional OS and user applications.

- This type of virtualization creates isolated containers on a single physical server and the OS instances to utilize the hard-ware and software in data centers.
- Containers behave like real servers. With containers you can create a portable, consistent operating environment for development, testing and deployment.
- This virtualization creates virtual hosting environments to allocates hardware resources among a large number of mutually distrusting users.
- Operating - system - level virtualization usually imposes little to no overhead, because programs in virtual partitions use the operating system's normal system call interface and do not need to be subjected to emulation or be run in an intermediate virtual machine.
- Operating system-level virtualization is not as flexible as other virtualization approaches since it cannot host a guest operating system different from the host one, or a different guest kernel.
- Instead of trying to run an entire guest OS, container virtualization isolates the guests, but doesn't try to virtualize the hardware. Instead, you have containers for each virtual environment.
- With container-based technologies, you'll need a patched kernel and user tools to run the virtual environments. The kernel provides process isolation and performs resource management.

Why operating system level virtualization is required ?

- Operating system level virtualization provides feasible solution for hardware level virtualization issue. It inserts a virtualization layer inside an operating system to partition a machine's physical resources.
- It enables multiple isolated VMs within a single operating system kernel. This kind of VM is often called a virtual execution environment (VE), Virtual Private System (VPS) or simply container.
- From the user's point of view, virtual execution environment look like real servers.
- This means a virtual execution environment has its own set of processes, file system, user accounts, network interfaces with IP addresses, routing tables, firewall rules etc.
- Although VEs can be customized for different people, they share the same operating system kernel. Therefore, OS-level virtualization is also called single-OS image virtualization.

Challenges to cloud computing in OS level virtualization ?

- Cloud computing is transforming the computing landscape by shifting the hardware and staffing costs of managing a computational center to third parties.
- Cloud computing has at least two challenges :
  1. The ability to use a variable number of physical machines and virtual machine instances depending on the needs of a problem. For example, a task may need only a single CPU during some phases of execution but may need hundreds of CPUs at other times.
  2. It is related to slow operation of instantiating new virtual machine. Currently, new virtual machines originate either as fresh boots or as replicates of a template VM, unaware of the current application state. Therefore, to better support cloud computing, a large amount of research and development should be done.

Advantages of OS virtualization :

1. OS virtualization provide least overhead among all types of virtualization solution.
2. They offer highest performance and highest density of virtual environment.
3. Low resource requirements.
4. High Scalability.

Disadvantage of OS virtualization :

1. They support only one operating system as base and guest OS in a single server.
2. It supports library level virtualization.

3.5 Application Virtualization

- Virtualization at the application level virtualizes an application as a VM. On a traditional OS, an application often runs as a process. Therefore, application-level virtualization is also known as process-level virtualization.
- A fully virtualized application is not installed in the traditional sense, although it is still executed as if it were. The application behaves at runtime like it is directly interfacing with the original operating system and all the resources managed by it, but can be isolated to varying degrees.
- Full application virtualization requires a virtualization layer. Application virtualization layers replace part of the runtime environment normally provided by the operating system.

- The layer intercepts all disk operations of virtualized applications and transparently redirects them to a virtualized location, often a single file.
- The application remains unaware that it accesses a virtual resource instead of a physical one. Since the application is now working with one file instead of many files spread throughout the system, it becomes easy to run the application on a different computer and previously incompatible applications can be run side-by-side.
- The most popular approach is to deploy High Level Language (HLL) VMs. Here the virtualization layer sits as an application program on top of the operating system, and the layer exports an abstraction of a VM that can run programs written and compiled to a particular abstract machine definition. Any program written in the HLL and compiled for this VM will be able to run on it.
- Benefits :
  1. Application virtualization uses fewer resources than a separate virtual machine.
  2. Application virtualization also enables simplified operating system migrations.
  3. Applications can be transferred to removable media or between computers without the need of installing them, becoming portable software.
- **Limitations :**
  1. Not all computer programs can be virtualized.
  2. Lower performance.

3.6 Virtual Clusters and Resource Management

- As with traditional physical servers, Virtual Machines (VMs) can also be clustered. A VM cluster starts with two or more physical servers.
- Most virtualization platforms, including XenServer and VMware ESX Server, support a bridging mode which allows all domains to appear on the network as individual hosts. By using this mode, VMs can communicate with one another freely through the virtual network interface card and configure the network automatically.
- Virtual clusters enable admins to deploy, track and manage containers across various systems to ensure performance, security and governance and low costs.
- With many VMs, an inefficient configuration always causes problems with overloading or underutilization.
- Amazon's EC2 provides elastic computing power in a cloud. EC2 permits customers to create VMs and to manage user accounts over the time of their use. Xen Server and VMware ESXi Server support a bridging mode which allows all domains to appear on the network as individual hosts. With this mode VMs can

communicate with one another freely through the virtual network interface card and configure the network automatically.

Physical versus virtual clusters :

- Virtual Clusters are built with VMs installed at one or more physical clusters. The VMs in a virtual cluster are interconnected by a virtual network across several physical networks.

Virtual cluster features :

a) Virtual machines can be restarted on other hosts if the host where the virtual machine running fails.

b) Distributed Resource Scheduler : Virtual machines can be load balanced so that none of the hosts is too overloaded or too much empty in the cluster.

c) Live migration : Of virtual machines from one host to other.

- Fig. 3.6.1 shows cloud platform with virtual cluster.

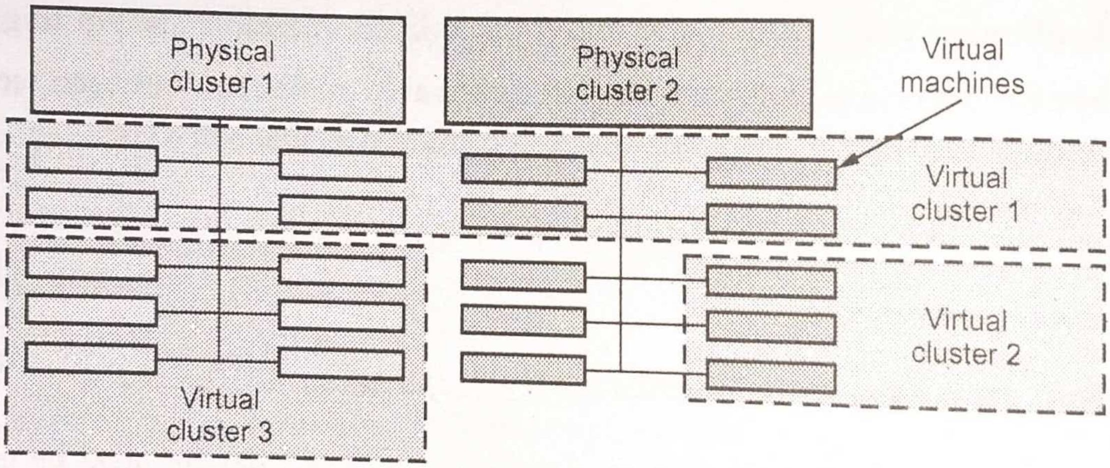

Fig. 3.6.1 Cloud platform example with three virtual clusters over two physical clusters

- The provisioning of VMs to a virtual cluster is done dynamically and they have the following properties :

  a) Virtual cluster nodes can be either physical or virtual with different operating systems.

  b) VM runs with a guest OS that manages the resources in the physical machine.

  c) The purpose of using VMs is to consolidate multiple functionalities on the same server.

  d) VMs can be replicated in multiple servers to promote parallelism, fault tolerance and disaster discovery.

  e) The no. of nodes in a virtual cluster can grow or shrink dynamically.

  f) The failure of some physical nodes will slow the work but the failure of VMs will cause no harm.

Characteristics virtual cluster :

1. Virtual machine or physical machine is used as virtual cluster nodes. Multiple VM running with different types of OS can be deployed on the same physical node.
2. Virtual machine runs with guest operating system. Host OS and VM OS are different but it manages the resources in the physical machine.
3. Virtual machine can be replicated in multiple servers and it support distributed parallelism, fault tolerance and disaster recovery.
4. Number of nodes of a virtual cluster may change accordingly.
5. If Virtual machine failes, it can not affect the host machine.

- **Virtual cluster is managed by four ways :**
  1. We can use a guest-based manager, by which the cluster manager resides inside a guest OS. Ex. : A Linux cluster can run different guest operating systems on top of the Xen hypervisor.
  2. We can bring out a host-based manager which itself is a cluster manager on the host systems. Ex. : VMware HA (High Availability) system that can restart a guest system after failure.
  3. An independent cluster manager, which can be used on both the host and the guest - making the infrastructure complex.
  4. Finally, we might also use an integrated cluster (manager), on the guest and host operating systems; here the manager must clearly distinguish between physical and virtual resources.

3.6.1 Virtualization in Disaster Recovery

- Data is a prime asset for all business organizations and it needs to be protected from getting lost, hacking, phishing and identity theft. Virtualization is the process of producing a virtual version of a system, software or even a working environment rather than a physical counterpart, as defined by the definition.
- Disaster Recovery (DR) relies upon the replication of data and computer processing in an off-premises location not affected by the disaster. When servers go down because of a natural disaster, equipment failure or cyber attack, a business needs to recover lost data from a second location where the data is backed up.
- With a disaster recovery plan, you can organize the actions to take in case of any disaster or incident. This will speed up the response time and minimize downtime.
- **Reducing downtime :** Virtualization software allows businesses to create image-based backups of their virtual machines. This means that in the event of a disaster, businesses can restore their systems quickly, rather than hours or days

needed to rebuild systems from scratch. Therefore, we can reduce and even eliminate downtime. Just access the data from another device to keep working.

- **Create off-site backups :** When we create business backups, we must consider having at least one copy of files on an off-site backup. This will allow us to rapidly recover files if anything helps with our local business data and storage device. When server can have issues, but power outages, fires and other natural disasters can also affect the place. Virtualization can automatically send backup files to off-site backup device.
- **Recover data from failed drives :** If hard drive dies or RAID fails, data virtualization system can help us. Since it is cloud storage, virtualization can keep copies of files even if storage devices stop working.
- **Test disaster recovery plans :** Virtualization can create a test environment on the system for testing the disaster recovery plan whenever required. This allows businesses to ensure that their disaster recovery plan is effective and can be executed efficiently when required.
- **Duplicate data for remote access :** With virtualization, businesses can duplicate their data in real-time or at specific intervals to a remote site. This allows remote users to access their data, applications and systems in the event of a disaster.

3.7 Introduction to Docker

- Docker is quickly changing the way that organizations are deploying software at scale.
- Docker is a tool that promises to easily encapsulate the process of creating a distributable artifact for any application, deploying it at scale into any environment, and streamlining the workflow and responsiveness of agile software organizations.
- **Benefits :**
  1. Packaging software in a way that leverages the skills developers already have.
  2. Bundling application software and required OS file systems together in a single standardized image format.
  3. Abstracting software applications from the hardware without sacrificing resources.

3.7.1 Process Simplification

- Docker can simplify both workflows and communication and that usually starts with the deployment story.
- Fig. 3.7.1 shows workflow with and without docker.

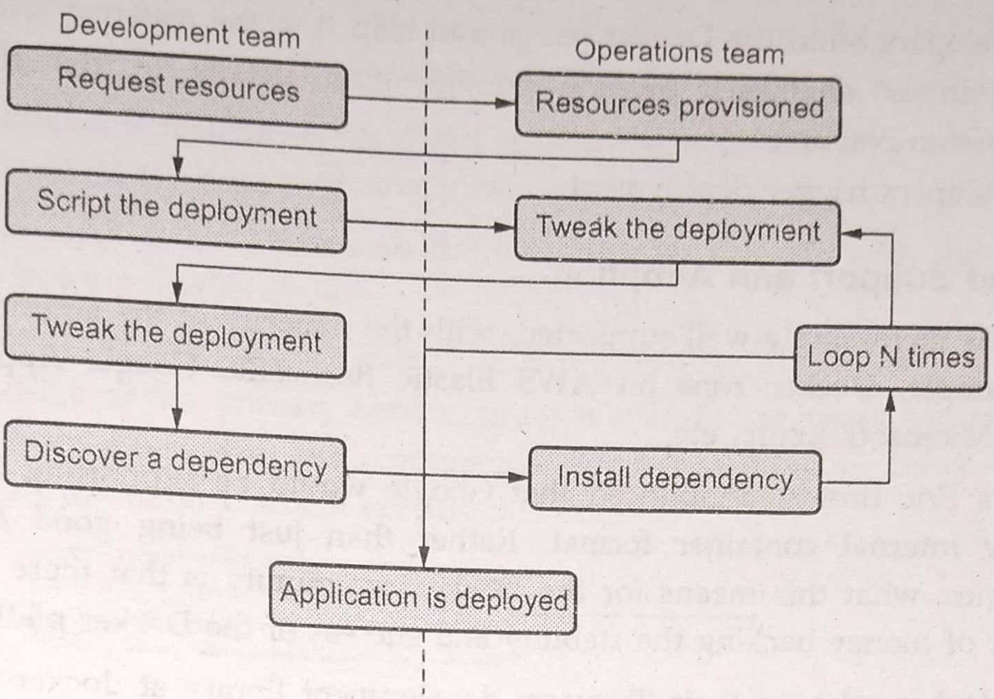

Fig. 3.7.1 Traditional deployment workflow (without docker)

1. Application developers request resources from operations engineers.
2. Resources are provisioned and handed over to developers.
3. Developers script and tool their deployment.
4. Operations engineers and developers tweak the deployment repeatedly.
5. Additional application dependencies are discovered by developers.
6. Operations engineers work to install the additional requirements.
7. Go to step 5 and 6.
8. The application is deployed.

- Fig. 3.7.2 shows Docker deployment workflow.

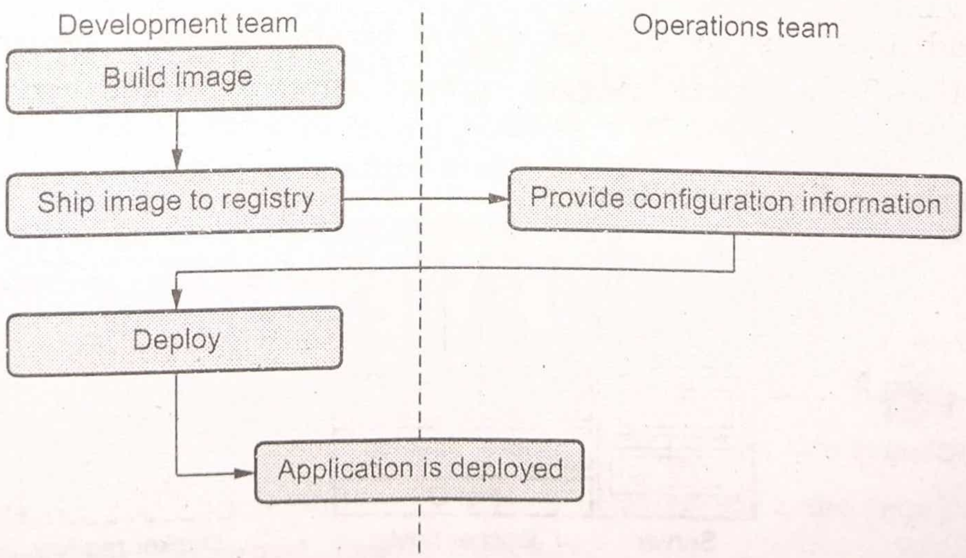

Fig. 3.7.2 Docker deployment workflow

1. Developers build the Docker image and ship it to the registry.
2. Operations engineers provide configuration details to the container and provision resources.
3. Developers trigger deployment.

3.7.2 Broad Support and Adoption

- Docker is increasingly well supported, with the majority of the large public clouds. For example, Docker runs on AWS Elastic Beanstalk, Google AppEngine, IBM Cloud, Microsoft Azure, etc.
- Google's Eric Brewer announced that Google would be supporting Docker as its primary internal container format. Rather than just being good PR for these companies, what this means for the Docker community is that there is starting to be a lot of money backing the stability and success of the Docker platform.
- When docker released their libswarm development library at docker-Con 2014, an engineer from Orchard demonstrated deploying a docker container to a heterogeneous mix of cloud providers at the same time.
- The Docker-client runs directly on most major operating systems, but because the Docker server uses Linux containers, it does not run on non-Linux systems.
- Docker has traditionally been developed on the Ubuntu Linux distribution, but today most Linux distributions and other major operating systems are now supported where possible.

3.7.3 Architecture

- The fundamental architecture of Docker is a simple client - server model, with only one executable that acts as both components, depending on how you invoke the docker command.
- Underneath those simple exteriors, Docker heavily leverages kernel mechanisms such as IPTABLES, virtual bridging, cgroups, namespaces and various filesystem drivers.
- Fig. 3.7.3 shows docker architecture.

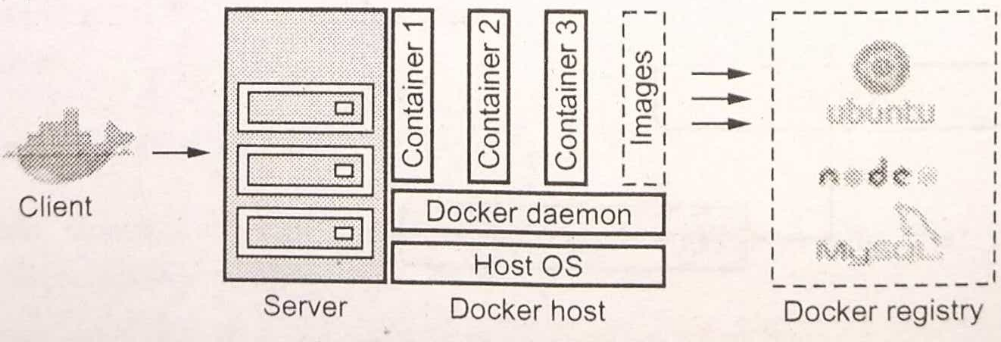

Fig. 3.7.3 Docker architecture

- It consists of two parts : The client and the server. Registry is one more components which stores docker images and metadata about those images.
- Docker Engine is a client-server based application with following components -
  1. A server which is a continuously running service called a **daemon process.**
  2. A REST API which interfaces the programs to use talk with the daemon and give instruct it what to do.
  3. A command line interface client.
- Docker client is the primary service using which docker users communicate with the docker. When we use commands "docker run" the client sends these commands to dockerd, which execute them out.

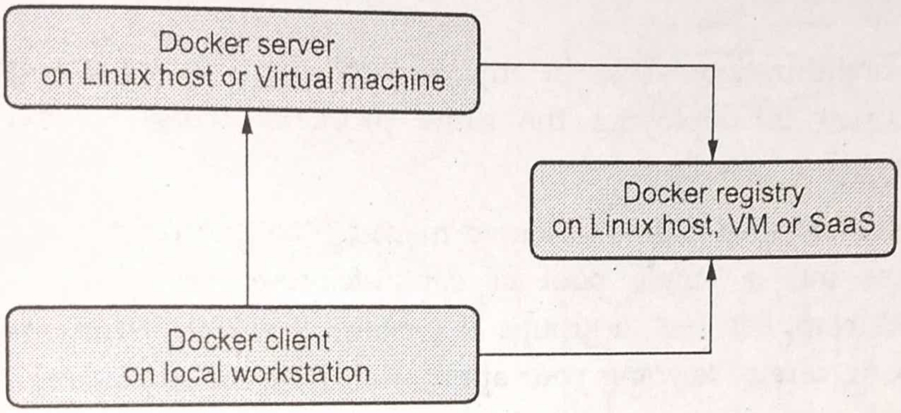

Fig. 3.7.4 Data flow

- The command used by docker depend on docker AP. In docker client can interact more than one daemon process.
- The docker images are building the block of docker or docker image is a read-only template with instructions to create a docker container. Docker images are the most build part of docker life cycle.
- The server does the ongoing work of running and managing your containers, and you use the client to tell the server what to do.
- The docker daemon can run on any number of servers in the infrastructure and a single client can address any number of servers.
- Clients drive all of the communication, but docker servers can talk directly to image registries when told to do so by the client.
- Clients are responsible for directing servers what to do and servers focus on hosting containerized applications.
- Docker registry keeps docker images. We can run our private registry.
- When we run the docker pull and docker run commands, the required images are pulled from our configured registry directory.

- Using docker push command, the image can be uploaded to our configured registry directory.

3.7.4 Container and Kubernetes

- A container image is a ready-to-run software package that includes everything a program needs to execute, including the code and any run-times it needs, application and system libraries and default values for any important settings.
- Container orchestration is concerned with the management of container lifecycles, particularly in large, dynamic environments. Container orchestration is used by software teams to control and automate a variety of tasks on container management.
- Container orchestration works in any context where containers are employed. It can assist you in deploying the same program across several environments without having to rewrite it.
- Kubernetes is an open-source container management platform that unifies a cluster of machines into a single pool of compute resources. With kubernetes, you organize your applications in groups of containers, which it runs using the Docker engine, taking care of keeping your application running as you request.
- Kubernetes is an open source container orchestration platform that automates many of the manual processes involved in deploying, managing, and scaling containerized applications.
- Kubernetes was originally developed and designed by engineers at Google.
- The primary responsibility of kubernetes is container orchestration. That means making sure that all the containers that execute various workloads are scheduled to run physical or virtual machines.
- The containers must be packed efficiently following the constraints of the deployment environment and the cluster configuration. In addition, kubernetes must keep an eye on all running containers and replace dead, unresponsive or otherwise unhealthy containers.
- Kubernetes uses docker to run images and manage containers.
- Kubernetes allows several containers to work in harmony, reducing operational burden. Interestingly, this includes docker containers. Kubernetes can be integrated with the docker engine and uses "Kubelets" to coordinate the scheduling of docker containers.
- The docker engine runs the container image, which is created by running docker build. The higher-level concepts (load balancing, service discovery and network policies) are controlled by kubernetes. When combined, both docker and

kubernetes can develop a modern cloud architecture. However, it should be remembered the two systems, at their core, are fundamentally different.

- Fig. 3.7.5 shows kubernetes architecture.

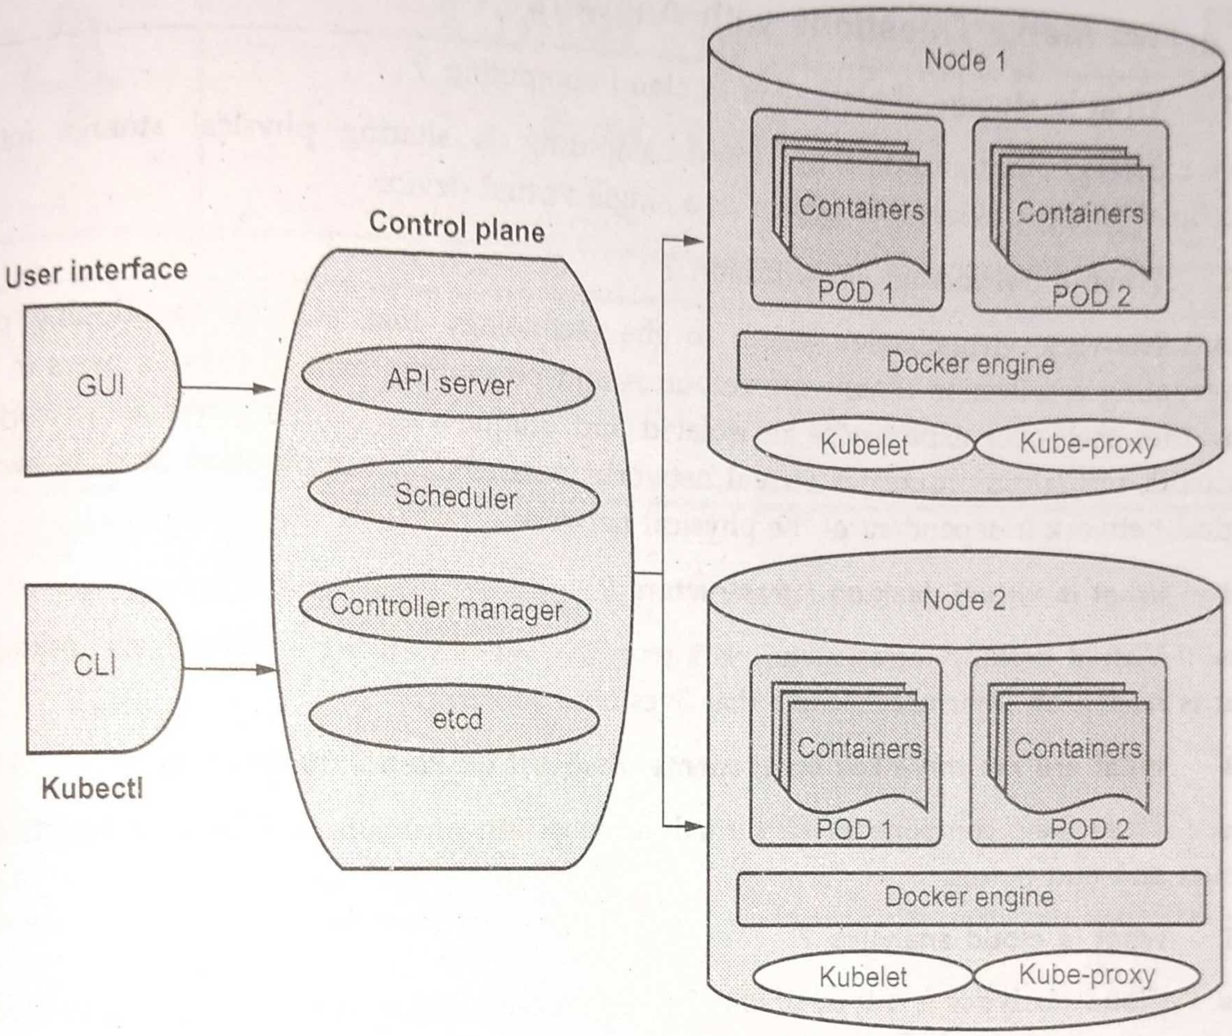

Fig. 3.7.5 Kubernetes architecture

- **Kubelet** : This function runs on nodes, reads container manifests, and assures defined containers have started and are running.
- **Node** : These perform the assigned tasks, with the kubernetes master controlling them.
- **Master** : This controls the kubernetes nodes and is the source of all task assignments.
- **Pod** : When one or more containers are deployed to one node. Containers in a pod will share a host name, an IP address, IPC and other resources.
- **Replication controller** : Controls the number of "identical" copies in a pod that should be running in different locations on the cluster.
- **Service** : This will decouple the work definitions from the pods. Service requests are automatically sent to the right pod, regardless of location.

- **Kubectl :** The primary configuration tool for kubernetes.
- **Kubernetes objects :** These are persistent entities within the Kubernetes system. They are used to represent the state of the cluster.

3.8 Two Marks Questions with Answers

Q.1 What is storage virtualization in cloud computing ?

Ans. : Storage virtualization in cloud computing is sharing physical storage into multiple storage devices that appear as a single virtual device.

Q.2 What is networking virtualization ?

Ans. : Network virtualization refers to the technology that enables partitioning or aggregating a collection of network resources and presenting them to various users in a way that each user experiences an isolated and unique view of the physical network. Network virtualization creates virtual networks whereby each application sees its own logical network independent of the physical network.

Q.3 What is virtual desktop infrastructure ?

Ans. : Virtual desktop infrastructure is a term that refers to using a virtualized desktop that is hosted on a virtual machine that lives on a server.

Q.4 What are the three key components of virtual desktop infrastructure ?

Ans. : Three key components of virtual desktop infrastructure are host, connection broker and end points.

Q.5 What is cloud analytics ?

Ans. : Cloud analytics is a type of cloud service model where data analysis and related services are performed on a public or private cloud. Cloud analytics can refer to any data analytics or business intelligence process that is carried out in collaboration with a cloud service provider.

Q.6 What is file level storage virtualization ?

Ans. : File level storage virtualization provides storage volumes to operating systems and applications in the form of files and directories. Access to storage is through network protocols, such as common Internet file system and network file systems. storage resources and capacity is may be underutilized because files are bound to a specific file server. It is necessary move the file from one server to another server.

□□□