

IT8001

Information Storage and Management

Professional Elective

Topic

Storage System
Environment



Varun C.M.

Assistant Professor, Department of Information Technology,
St. Xavier's Catholic College of Engineering

CO1

Understand the logical and physical components of a Storage infrastructure.

Application

The application sends requests to the underlying operating system to perform read/write (R/W) operations on the storage devices.

Applications can be layered on the database, which in turn uses the OS services to perform R/W operations on the storage devices.

Some examples of these applications are e-mail, enterprise resource planning (ERP), decision support system (DSS), resource management, backup, authentication and antivirus applications, and so on.

Database Management System (DBMS)

A DBMS controls the creation, maintenance, and use of a database.

The DBMS processes an application's request for data and instructs the operating system to transfer the appropriate data from the storage.

Host (Compute)

The computers on which the applications run are referred to as hosts or compute systems.

Hosts can be physical or virtual machines.

Physical Host - desktop computers, servers or a cluster of servers, laptops, and mobile devices.

Compute Virtualization software enables creating **virtual machines** on top of a physical compute infrastructure.

Host (Compute)

1. Operating System

An operating system controls all aspects of computing.

It works between the application and the physical components of a compute system.

An operating system also performs basic storage management tasks while managing other underlying components, such as the file system, volume manager, and device drivers.

In a virtualized compute environment, the virtualization layer works between the operating system and the hardware resources. Here the OS might work differently based on the type of compute virtualization implemented.

Host (Compute)

1. **Operating System** - *contd...*

Memory Virtualization is an operating system feature that virtualizes the physical memory (RAM) of a host. It creates virtual memory with an address space larger than the physical memory space present in the compute system.

The operating system utility that manages the virtual memory is known as the virtual memory manager (VMM). The VMM manages the virtual-to-physical memory mapping and fetches data from the disk storage when a process references a virtual address that points to data at the disk storage.

The space used by the VMM on the disk is known as a swap space. A swap space (also known as page file or swap file) is a portion of the disk drive that appears to be physical memory to the operating system.

Host (Compute)

1. **Operating System** - contd...

A process known as paging moves inactive physical memory pages onto the swap file and brings them back to the physical memory when required.

This enables efficient use of the available physical memory among different applications.

The operating system typically moves the least used pages into the swap file so that enough RAM is available for processes that are more active.

Host (Compute)

2. Device Driver

A device driver is special software that permits the operating system to interact with a specific device, such as a printer, a mouse, or a disk drive.

A device driver enables the operating system to recognize the device and to access and control devices.

Device drivers are hardware-dependent and operating-system-specific.

Host (Compute)

3. Volume Manager

The evolution of Logical Volume Managers (LVMs) enabled dynamic extension of file system capacity and efficient storage management.

The LVM is software that runs on the compute system and manages logical and physical storage.

LVM is an intermediate layer between the file system and the physical disk.

It can partition a larger-capacity disk into virtual, smaller-capacity volumes or aggregate several smaller disks to form a larger virtual volume.

Host (Compute)

3. **Volume Manager** - contd...

Disk partitioning was introduced to improve the flexibility and utilization of disk drives.

In partitioning, a disk drive is divided into logical containers called logical volumes (LVs)

Concatenation is the process of grouping several physical drives and presenting them to the host as one big logical volume

Host (Compute)

3. **Volume Manager** - contd...

A volume group is created by grouping together one or more physical volumes.

A unique physical volume identifier (PVID) is assigned to each physical volume when it is initialized for use by the LVM.

Physical volumes can be added or removed from a volume group dynamically.

Host (Compute)

4. File System

A file is a collection of related records or data stored as a unit with a name. A file system is a hierarchical structure of files.

A file system enables easy access to data files residing within a disk drive, a disk partition, or a logical volume.

A file system consists of logical structures and software routines that control access to files. It provides users with the functionality to create, modify, delete, and access files.

Host (Compute)

4. **File System** - *contd...*

Examples of common file systems are

- **FAT 32** (File Allocation Table) for Microsoft Windows
- **NTFS** (NT File System) for Microsoft Windows
- **UFS** (UNIX File System) for UNIX
- **EXT2/3** (Extended File System) for Linux

Host (Compute)

4. **File System** - *contd...*

The file system tree starts with the root directory. The root directory has a number of subdirectories.

A file system should be mounted before it can be used.

A file system can be either a journaling file system or a non-journaling file system.

Host (Compute)

4. **File System** - contd...

Non-journaling file systems cause a potential loss of files because they use separate writes to update their data and metadata. If the system crashes during the write process, the metadata or data might be lost or corrupted.

A **journaling** file system uses a separate area called a log or journal. This journal might contain all the data to be written (physical journal) or just the metadata to be updated (logical journal).

Before changes are made to the file system, they are written to this separate area. After the journal has been updated, the operation on the file system can be performed.

If the system crashes during the operation, there is enough information in the log to “replay” the log record and complete the operation.

Host (Compute)

5. Compute Virtualization

Compute virtualization is a technique for masking or abstracting the physical hardware from the operating system.

It enables multiple operating systems to run concurrently on single or clustered physical machines. This technique enables creating portable virtual compute systems called virtual machines (VMs).

Each VM runs an operating system and application instance in an isolated manner.

Compute virtualization is achieved by a virtualization layer that resides between the hardware and virtual machines. This layer is also called the hypervisor

Host (Compute)

5. Compute Virtualization - contd...

The hypervisor provides hardware resources, such as CPU, memory, and network to all the virtual machines. Within a physical server, a large number of virtual machines can be created depending on the hardware capabilities of the physical server.

Compute virtualization allows multiple operating systems and applications to run on a single physical machine. This technique significantly improves server utilization and provides server consolidation.

Connectivity

Connectivity refers to the interconnection between hosts or between a host and peripheral devices, such as printers or storage devices.

Physical Components of Connectivity

The physical components of connectivity are the hardware elements that connect the host to storage.

Three physical components of connectivity between the host and storage are the host interface device, port, and cable.

Connectivity

A host interface device or host adapter connects a host to other hosts and storage devices.

Examples of host interface devices are host bus adapter (HBA) and network interface card (NIC).

Host bus adaptor is an application-specific integrated circuit (ASIC) board that performs I/O interface functions between the host and storage, relieving the CPU from additional I/O processing workload. A host typically contains multiple HBAs.

A port is a specialized outlet that enables connectivity between the host and external devices.

Connectivity

Interface Protocols

A protocol enables communication between the host and storage. Protocols are implemented using interface devices (or controllers) at both source and destination.

The popular interface protocols used for host to storage communications are Integrated Device Electronics/Advanced Technology Attachment (IDE/ATA), Small Computer System Interface (SCSI), Fibre Channel (FC) and Internet Protocol (IP).

Connectivity

Interface Protocols

1. *IDE/ATA and Serial ATA*

IDE/ATA is a popular interface protocol standard used for connecting storage devices, such as disk drives and CD-ROM drives. This protocol supports parallel transmission and therefore is also known as Parallel ATA (PATA) or simply ATA.

Connectivity

Interface Protocols

2. *SCSI and Serial SCSI*

SCSI has emerged as a preferred connectivity protocol in high-end computers. This protocol supports parallel transmission and offers improved performance, scalability, and compatibility compared to ATA.

However, the high cost associated with SCSI limits its popularity among home or personal desktop users.

Connectivity

Interface Protocols

3. *Fibre Channel*

Fibre Channel is a widely used protocol for high-speed communication to the storage device.

The Fibre Channel interface provides gigabit network speed. It provides a serial data transmission that operates over copper wire and optical fiber.

Connectivity

Interface Protocols

4. Internet Protocol


IP is a network protocol that has been traditionally used for host-to-host traffic. With the emergence of new technologies, an IP network has become a viable option for host-to-storage communication.

IP offers several advantages in terms of cost and maturity and enables organizations to leverage their existing IP-based network.

Storage

Storage is a core component in a data center.

A storage device uses magnetic, optic, or solid state media. Disks, tapes, and diskettes use magnetic media, whereas CD/DVD uses optical media for storage. Removable Flash memory or Flash drives are examples of solid state media.



References

EMC Corporation, "Information Storage and Management" , Wiley, India.
Robert Spalding, "Storage Networks: The Complete Reference", Tata McGraw Hill, Osborne, 2003.



Varun C.M.

Assistant Professor, Department of Information Technology, St. Xavier's Catholic College of Engineering