USN | | | | | | | | | |

Internal Assessment Test 1 – Nov 2021

Sub:	Storage Area Networks				Sub Code:	15/17CS753	Branch:	CSE/ISE	
Date:	13/11/2021	Duration:	90 mins	Max Marks: 50	Sem / Sec:	VII/D			OBE

Answer any FIVE FULL Questions

		MARKS	CO	RBT
1	Define Data center? Explain the Key Characteristics of Data Center.	[10]	CO1	L2

Diagram carries 3 M

1. **Availability**

- In data-center, all core-elements must be designed to ensure availability (Figure 1-5).
- If the users cannot access the data in time, then it will have negative impact on the company. (For example, if amazon server goes down for even 5 min, it incurs huge loss in millions).

2. Security

- To prevent unauthorized-access to data,
  1. Good polices & procedures must be used.
  2. Proper integration of core-elements must be established.
- Security-mechanisms must enable servers to access only their allocated-resources on the storage.

3. Scalability

- It must be possible to allocate additional resources on-demand w/o interrupting normal-operations.
- The additional resources includes CPU-power and storage.
- Business growth often requires deploying
  → more servers
  → new applications and
  → additional databases.
- The storage-solution should be able to grow with the company.

4. Performance

- All core-elements must be able to
  → provide optimal-performance and
  → service all processing-requests at high speed.
- The data-center must be able to support performance-requirements.

5. Data Integrity

- Data integrity ensures that data is written to disk exactly as it was received.
- For example: Parity-bit or ECC (error correction code).

- Data-corruption may affect the operations of the company.

6. Storage Capacity

- The data-center must have sufficient resources to store and process large amount of data efficiently.
- When capacity-requirement increases, the data-center must be able
  → to provide additional capacity without interrupting normal-operations.
- Capacity must be managed by reallocation of existing-resources rather than by adding new resources

7. Manageability

- A data-center must perform all operations and activities in the most efficient manner.
- Manageability is achieved through automation i.e. reduction of human-intervention in common tasks.

1. Managing a Data Center

- Managing a data-center involves many tasks.
- Key management-tasks are: 1) Monitoring 2) Reporting and 3) Provisioning.

1. **Monitoring** is a process of continuous
   → collection of information and
   → review of the entire storage infrastructure (called as Information Storage System).

- Following parameters are monitored:
  i. Security
  ii. Performance
  iii. Accessibility and
  iv. Capacity.

2. **Reporting** is done periodically on performance, capacity and utilization of the resources.

- Reporting tasks help to
  → establish business-justifications and
  → establish chargeback of costs associated with operations of data-center.

3. **Provisioning** is process of providing h/w, s/w & other resources needed to run a data-center.

- Main tasks are: i) Capacity Planning and ii) Resource Planning.

7 Points each carriers (1*7 =7Marks)

Capacity Planning

- It ensures that future needs of both user & application will be addressed in most cost-effective way

ii. Resource Planning

- It is the process of evaluating & identifying required resources such as
  → Personnel (employees)
  → Facility (site or plant) and
  → Technology (Artificial Intelligence, Deep Learning).

&

2	Briefly discuss the different components used in Data Center.	[10]	CO1	L2

- The data flows from an application to storage through various components collectively referred as a data-center environment.
- The five main components in this environment are
  1. Application
  2. DBMS

3. Host
4. Connectivity and
5. Storage.

- These entities, along with their physical and logical-components, facilitate data-access.

Each Components carriers 2 Marks (5*2=10)

Application

- An application is a program that provides the logic for computing-operations.
- It provides an interface between user and host. (R/W --> read/write)
- The application sends requests to OS to perform R/W-operations on the storage.
- Applications can be placed on the database.
  Then, the database can use OS-services to perform R/W-operations on the storage.
- Applications can be classified as follows:
  → business applications
  → infrastructure management applications
  → data protection applications
  → security applications.
- Some examples of the applications are:
  → e-mail
  → enterprise resource planning (ERP)
  → backup
  → antivirus
- Characteristics of I/Os generated by application influence the overall performance of storage-device.
- Common I/O characteristics are:
  → Read vs. Write intensive
  → Sequential vs. Random
  → I/O size

1. DBMS

- DBMS is a structured way to store data in logically organized tables that are inter-related.
- The DBMS
  → processes an application's request for data and
  → instructs the OS to transfer the appropriate data from the storage.
- Advantages:
  1. Helps to optimize the storage and retrieval of data.
  2. Controls the creation, maintenance and use of a database.

2. Host

- Host is a client- or server-computer that runs applications.
- Users store and retrieve data through applications. (hosts --> compute-systems)
- Hosts can be physical- or virtual-machines.
- Example of host includes
  → desktop computers
  → servers
  → laptops
  → smartphones.

- A host consists of

1. CPU
2. Memory
3. I/O devices
4. Software. The software includes

i. OS
ii. Device-drivers
iii. Logical volume manager (LVM)
iv. File-system

- The software can be installed individually or may be part of the OS.

Operating System (OS)

- An OS is a program that acts as an intermediary between
  → application and
  → hardware-components.
- The OS controls all aspects of the computing-environment.
- Data-access is one of the main service provided by OS to the application.
- Tasks of OS:
  1. Monitor and respond to user actions and the environment.
  2. Organize and control hardware-components.
  3. Manage the allocation of hardware-resource (simply the resource).
  4. Provide security for the access and usage of all managed resources.
  5. Perform storage-management tasks.
  6. Manage components such as file-system, LVM & device drivers.

Memory Virtualization

- Memory-virtualization is used to virtualize the physical-memory (RAM) of a host.
- It creates a VM with an address-space larger than the physical-memory space present in computer.
- The virtual-memory consists of
  → address-space of the physical-memory and
  → part of address-space of the disk-storage.
- The entity 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
- The space used by the VMM on the disk is known as a swap-space.
- A **swap-space** is a portion of the disk that appears like physical-memory to the OS.
- The memory is divided into contiguous blocks of fixed-size pages. (VM --> virtual-memory)

  Paging
  - A paging
    → moves inactive-pages onto the swap-file and
    → brings inactive-pages back to the physical-memory when required.
  - Advantages:
    1. Enables efficient use of the available physical-memory among different applications.

¤ Normally, the OS moves the least used pages into the swap-file.
¤ Thus, sufficient RAM is provided for processes that are more active.

- Disadvantage:
  1) Access to swap-file pages is slower than physical-memory pages. This is because
  → swap-file pages are allocated on the disk which is slower than physical-memory.

Device Driver

- It is a special software that permits the OS & hardware-component to interact with each other.
- The hardware-component includes printer, a mouse and a hard-drive.
- A device-driver enables the OS to
  → recognize the device and
  → use a standard interface to access and control devices.
- Device-drivers are hardware-dependent and OS-specific.

Logical Volume Manager (LVM)

- LVM is a software that
  → runs on the host and
  → manages the logical- and physical-storage.
- It is an intermediate-layer between file-system and disk.
- Advantages:
  1. Provides optimized storage-access.
  2. Simplifies storage-management. (PVID --> Physical-Volume IDentifier)
  3. Hides details about disk and location of data on the disk.
  4. Enables admins to change the storage-allocation without interrupting normal-operations.
  5. Enables dynamic-extension of storage-capacity of the file-system.
- The main components of LVM are: 1) Physical-volumes 2) Volume-groups and 3) Logical-volumes.
  1. **Physical-Volume (PV):** refers to a disk connected to the host.
  2. **Volume-Group (VG):** refers to a group of one or more PVs.
  - A unique PVID is assigned to each PV when it is initialized for use.
  - PVs can be added or removed from a volume-group dynamically.
  - PVs cannot be shared between different volume-groups.
  - The volume-group is handled as a single unit by the LVM.
  - Each PV is divided into equal-sized data-blocks called **physical-extents**.

  3. **Logical-Volume (LV):** refers to a partition within a volume-group.
  - Logical-volumes vs. Volume-group

  LV can be thought of as a disk-partition.
  Volume-group can be thought of as a disk.
  - The size of a LV is based on a multiple of the physical-extents.
  - The LV appears as a physical-device to the OS.
  - A LV is made up of non-contiguous physical-extents and may span over multiple PVs.
  - A file-system is created on a LV.
  - These LVs are then assigned to the application.

- A LV can also be mirrored to improve data-availability.

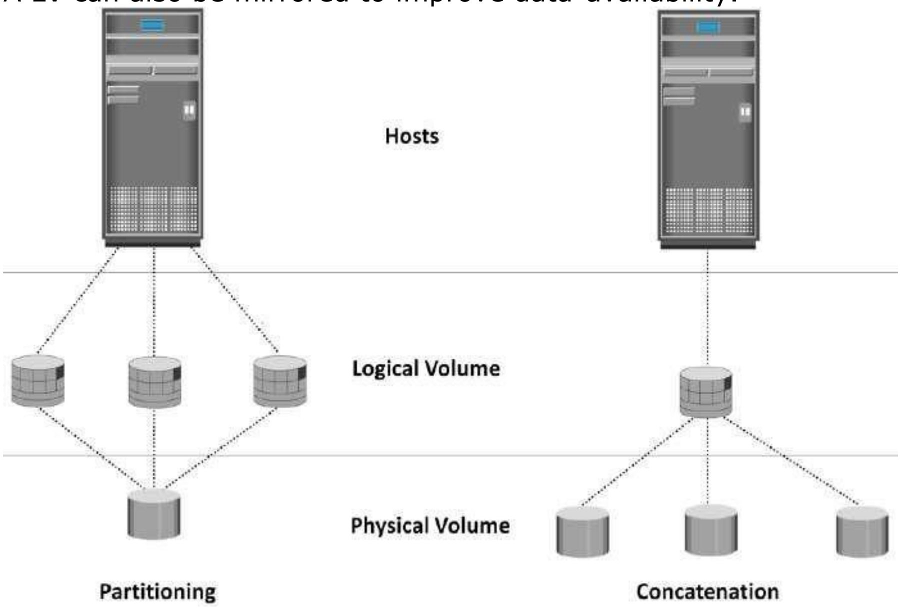

Figure 1-6: Disk partitioning and concatenation

- It can perform partitioning and concatenation (Figure 1-6).

1. Partitioning

- A larger-capacity disk is partitioned into smaller-capacity virtual-disks.
- Disk-partitioning is used to improve the utilization of disks.

2. Concatenation

- Several smaller-capacity disks are aggregated to form a larger-capacity virtual-disk.
- The larger-capacity virtual-disk is presented to the host as one big logical-volume.

3 Consider a Disk I/O System in which an I/O request arrives at the rate of 80 IOPS. Service time is 3 ms. Compute the following (i) Utilization (ii) Total Response Time. (iii) Average Queue Size (iv) Time Spent by a request in a queue. [10] CO1 L3

4 With a neat diagram Explain the different RAID Level and Techniques [10] CO1 L3

3 RAID Techniques carries 1 mark each (3*1=3 Marks)

- RAID-levels are defined based on following 3 techniques:
  1. Striping (used to improve performance of storage)
  2. Mirroring (used to improve data-availability) and
  3. Parity (used to provide data-protection)
- The above techniques determine

  → performance of storage-device (i.e. better performance --> least response-time)

  → data-availability

  → data-protection

- Some RAID-arrays use a combination of above 3 techniques. For example:
  Striping with mirroring
  Striping with parity

1. Striping

- Striping is used to improve performance of a storage-device.
- It is a technique of splitting and distribution of data across multiple disks.
- Main purpose: To use the disks in parallel.
- It can be bitwise, byte-wise or block wise.
- A **RAID-set** is a group of disks.
-

Figure 1-11: Striped RAID set

- In each disk, a predefined number of strips are defined.
- **Strip** refer to a group of continuously-addressable-blocks in a disk.
- **Stripe** refer to a set of aligned-strips that spans all the disks. (Figure 1-11)
- **Strip-size** refers to maximum amount-of-data that can be accessed from a single disk. In other words, strip-size defines the number of blocks in a strip.
- In a stripe, all strips have the same number of blocks.
- **Stripe-width** refers to the number of strips in a stripe.
- Striped-RAID does not protect data. To protect data, parity or mirroring must be used.
- Advantage:
  1. As number of disks increases, the performance also increases. This is because
     → more data can be accessed simultaneously.

(Example for stripping: If one man is asked to write A-Z the amount of time taken by him will be more as compared to 2 men writing A-Z because from the 2 men, one man will write A-M and another will write N-Z at the same time so this will speed up the process)

1. Mirroring

- Mirroring is used to improve data-availability (or data-redundancy).
- All the data is written to 2 disks simultaneously. Hence, we have 2 copies of the data.
- Advantages:

1. Reliable

- Provides protection against single disk-failure.
- In case of failure of one disk, the data can be accessed on the surviving-disk (Figure 1-12).
- Thus, the controller can still continue to service the host's requests from surviving-disk.
- When failed-disk is replaced with a new-disk, controller copies data from surviving-disk to new-disk
- The disk-replacement activity is transparent to the host.

2. Increases read-performance because each read-request can be serviced by both disks.

- Disadvantages:
  1. Decreases write-performance because
     → each write-request must perform 2 write-operations on the disks.
  2. Duplication of data. Thus, amount of storage-capacity needed is twice amount of data being stored (E.g. To store 100GB data, 200GB disk is needed).

3. Considered expensive and preferred for mission-critical applications (like military application).

- Mirroring is not a substitute for data-backup. Mirroring vs. Backup
  1. Mirroring constantly captures changes in the data.
  2. On the other hand, backup captures point-in-time images of data.

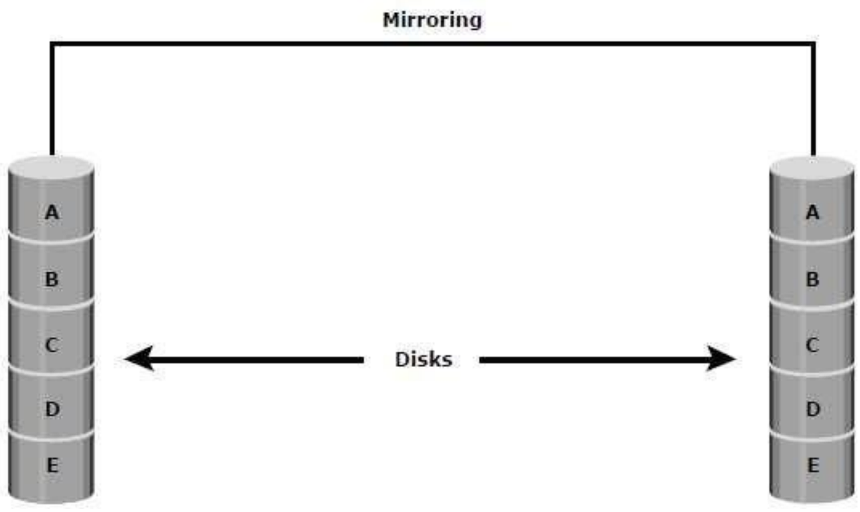

Figure 1-12: Mirrored disks in an array

2. **Parity**

- Parity is used to provide data-protection in case of a disk-failure.
- An additional disk is added to the stripe-width to hold parity.
- In case of disk-failure, parity can be used for reconstruction of the missing-data.
- Parity is a technique that ensures protection of data without maintaining a duplicate-data
- Parity-information can be stored on
  → separate, dedicated-disk or
  → distributed across all the disks.
- For example (Figure 1-13):
  Consider a RAID-implementation with 5 disks (5*100 GB = 500 GB).
  1. The first four disks contain the data (4*100 = 400GB).
  2. The fifth disk stores the parity-information (1*100 = 100GB).

Parity vs. Mirroring

Parity requires 25% extra disk-space. (i.e. 500GB disk for 400GB data).
Mirroring requires 100% extra disk-space.(i.e. 800GB disk for 400GB data).

- The controller is responsible for calculation of parity.

- Parity-value can be calculated by

  P = D1 + D2 + D3 + D4

  where D1 to D4 is striped-data across the set of five disks.
- Now, if one of the disks fails (say D1), the missing-value can be calculated by D1 = P - (D2 + D3 + D4)

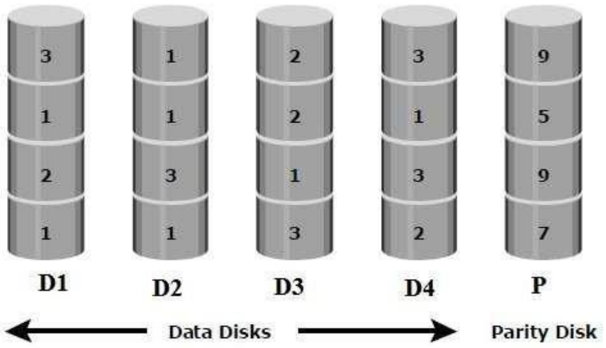

Figure 1-13: Parity RAID

- Advantages:
  1. Compared to mirroring, parity reduces the cost associated with data-protection.
  2. Compared to mirroring, parity consumes less disk-space. In previous example,
     i. Parity requires 25% extra disk-space. (i.e. 500GB disk for 400GB data).
     ii. Mirroring requires 100% extra disk-space. (i.e. 800GB disk for 400GB data).
- Disadvantage:
  1) Decreases performance of storage-device. For example:
     ¤ Parity-information is generated from data on the disk.
     ¤ Therefore, parity must be re-calculated whenever there is change in data.
     ¤ This re-calculation is time-consuming and hence decreases the performance.

RAID Levels

LEVELS	BRIEF DESCRIPTION
RAID 0	Striped array with no fault tolerance
RAID 1	Disk mirroring
RAID 3	Parallel access array with dedicated parity disk
RAID 4	Striped array with independent disks and a dedicated parity disk
RAID 5	Striped array with independent disks and distributed parity
RAID 6	Striped array with independent disks and dual distributed parity
Nested	Combinations of RAID levels. Example: RAID 1 + RAID 0

Table 1-1: Raid Levels

1. **RAID-0**

- RAID-0 is based on striping-technique (Figure 1-14).
- Striping is used to improve performance of a storage-device.
- It is a technique of splitting and distribution of data across multiple disks.
- Main purpose:
  To use the disks in parallel.
- Therefore, it utilizes the full storage-capacity of the storage-device.
- Read operation: To read data, all the strips are combined together by the controller.
- Advantages:
  1. Used in applications that need high I/O-throughput. (Throughput --> Efficiency).
  2. As number of disks increases, the performance also increases. This is because
     → more data can be accessed simultaneously.
- Disadvantage:
  1) Does not provide data-protection and data-availability in case of disk-failure.

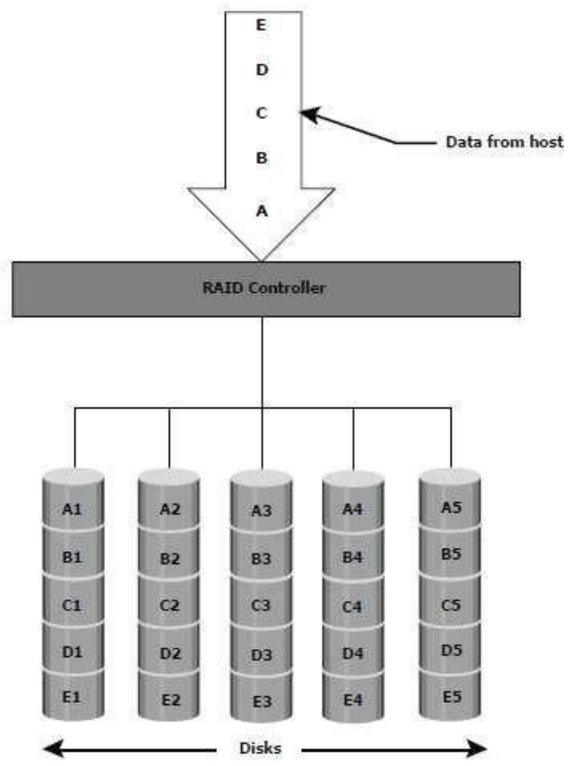

Figure 1-14: RAID 0

1. **RAID-1**

- RAID-1 is based on mirroring-technique.
- Mirroring is used to improve data-availability (or data-redundancy).
- Write operation: The data is stored on 2 different disks. Hence, we have 2 copies of data.
- Advantages:
  1. Reliable
  - Provides protection against single disk-failure.
  - In case of failure of one disk, the data can be accessed on the surviving-disk (Figure 1-15).

- Thus, the controller can still continue to service the host's requests from surviving-disk.
- When failed-disk is replaced with a new-disk, controller copies data from surviving-disk to new-disk
- The disk-replacement activity is transparent to the host.

2. Increases read-performance because each read-request can be serviced by both disks.

- Disadvantages:
  1. Decreases write-performance because
     → each write-request must perform 2 write-operations on the disks.

2. Duplication of data.
Thus, amount of storage-capacity needed is twice amount of data being stored (E.g. To store 100 GB data, 200 GB disk is required).
3. Considered expensive and preferred for mission-critical applications (like military application)

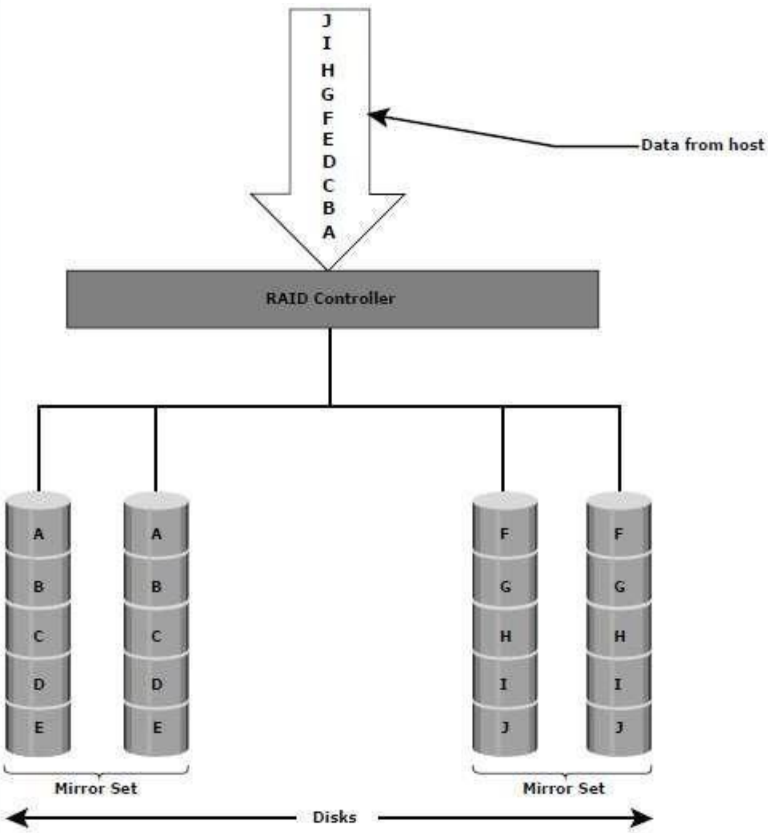

Figure 1-15: RAID-1

2. **Nested-RAID**

Figure 1-16: Nested-RAID

- Most data-centers require data-availability & performance from their storage-devices (Figure 1-16).
- RAID-01 and RAID-10 combines
  → performance-benefit of RAID-0 and
  → availability-benefit of RAID-1. (RAID-10 is also known as RAID-1+0)

- It uses mirroring- and striping-techniques.
- It requires an even-number of disks. Minimum no. of disks = 4.
- Some applications of RAID-10:
    1. High transaction rate OLTP (Online Transaction Processing)
    2. Large messaging installations
    3. Database applications that require
        → high I/O-throughput
        → random-access and
        → high-availability.
- Common misunderstanding is that RAID-10 and RAID-01 are the same. But, they are totally different

1. RAID-10

- RAID-10 is also called **striped-mirror**.
- The basic element of RAID-10 is a mirrored-pair.
    1. Firstly, the data is mirrored and
    2. Then, both copies of data are striped across multiple-disks.

2. RAID-01

- RAID-01 is also called **mirrored-stripe**
- The basic element of RAID-01 is a stripe.
    1. Firstly, data are striped across multiple-disks and
    2. Then, the entire stripe is mirrored.

- Advantage of rebuild-operation::

1. Provides protection against single disk-failure.

- In case of failure of one disk, the data can be accessed on the surviving-disk (Figure 1-15).
- Thus, the controller can still continue to service the host's requests from surviving-disk.
- When failed-disk is replaced with a new-disk, controller copies data from surviving-disk to new-disk

- Disadvantages of rebuild-operation:
    1. Increased and unnecessary load on the surviving-disks.
    2. More vulnerable to a second disk-failure.

3. RAID-3

- RAID-3 uses both striping & parity techniques.
    1. Striping is used to improve performance of a storage-device.
    2. Parity is used to provide data-protection in case of disk-failure.
- Parity-information is stored on separate, dedicated-disk.
- Data is striped across all disks except the parity-disk in the array.
- In case of disk-failure, parity can be used for reconstruction of the missing-data.
- For example (Figure 1-17):
    - Consider a RAID-implementation with 5 disks (5*100GB = 500GB).
        1. The first 4 disks contain the data (4*100 = 400GB).
        2. The fifth disk stores the parity-information (1*100 = 100GB).
    - Therefore, parity requires 25% extra disk-space (i.e. 500GB disk for 400GB data).
- Advantages:
    1. Striping is done at the bit-level.
        Thus, RAID-3 provides good bandwidth for the transfer of large volumes of data.

2. Suitable for video streaming applications that involve large sequential data-access.

- Disadvantages:
    1. Always reads & writes complete stripes of data across all disks '.' disks operate in parallel.
    2. There are no partial writes that update one out of many strips in a stripe.

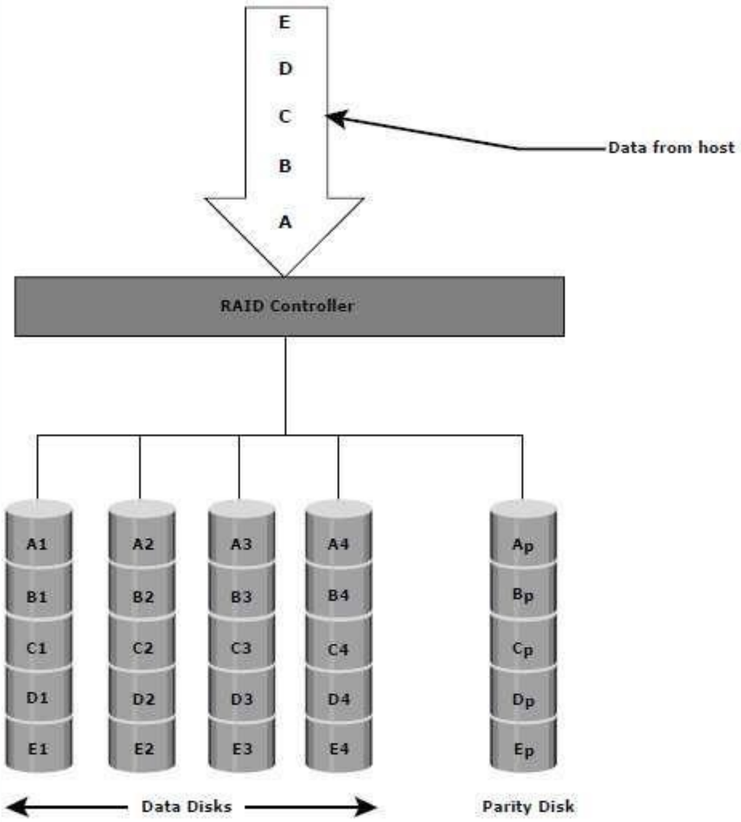

Figure 1-17: RAID-3

4. **RAID-4**

- Similar to RAID-3, RAID-4 uses both striping & parity techniques.
    1. Striping is used to improve performance of a storage-device.
    2. Parity is used to provide data-protection in case of disk-failure.
- Parity-information is stored on a separate dedicated-disk.
- Data is striped across all disks except the parity-disk.
- In case of disk-failure, parity can be used for reconstruction of the missing-data.
- Advantages:
    1. Striping is done at the block-level.
       Hence, data-element can be accessed independently.
       i.e. A specific data-element can be read on single disk without reading an entire stripe

2. Provides
   → good read-throughput and
   → reasonable write-throughput.

5. **RAID-5**

- Problem:

In RAID-3 and RAID-4, parity is written to a dedicated-disk.

If parity-disk fails, we will lose our entire backup.

Solution: To overcome this problem, RAID-5 is proposed.

In RAID-5, we distribute the parity-information evenly among all the disks.

- RAID-5 similar to RAID-4 because
  it uses striping and the drives (strips) are independently accessible.
- Advantages:
  1. Preferred for messaging & media-serving applications.
  2. Preferred for RDBMS implementations in which database-admins can optimize data-access.

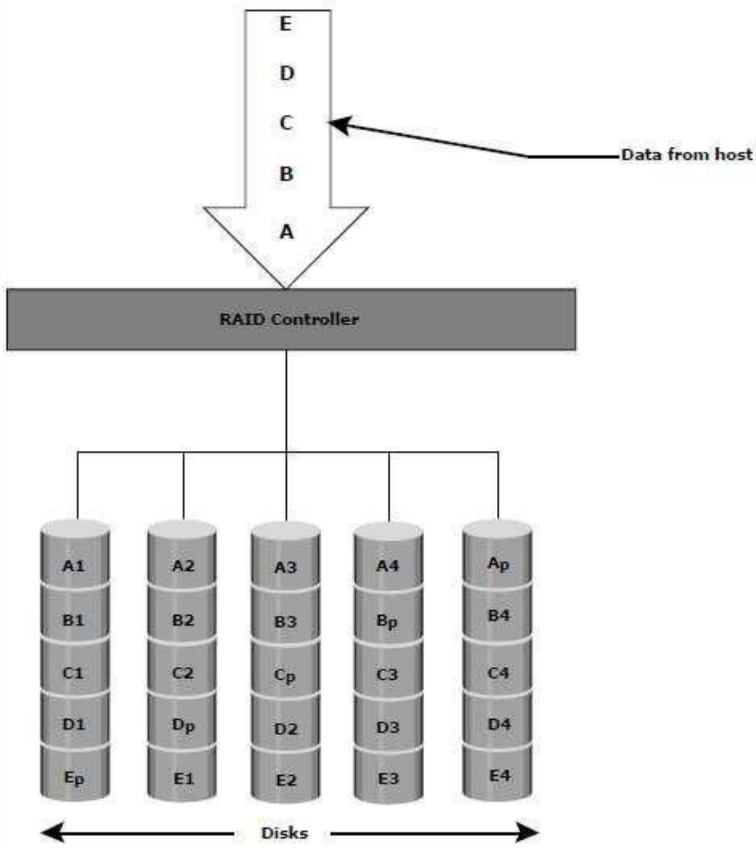

Figure 1-18: RAID-5

6. **RAID-6**

- RAID-6 is similar to RAID-5 except that it has
  → a second parity-element to enable survival in case of 2 disk-failures. (Figure 1-19).
- Therefore, a RAID-6 implementation requires at least 4 disks.
- Similar to RAID-5, parity is distributed across all disks.
- Disadvantages:
  Compared to RAID-5,
  1. Write-penalty is more. .'. RAID-5 writes perform better than RAID-6
  2. The rebuild-operation may take longer time. This is due to the presence of 2 parity-sets.

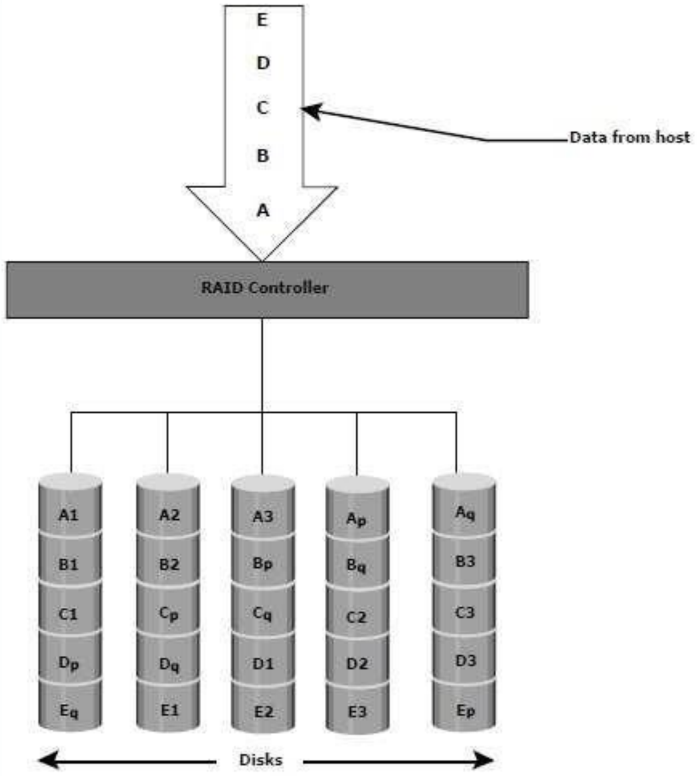

Figure 1-19: RAID-6

7 RAID level with diagram carries 1 mark each (7*1=7 Marks)

5	Briefly discuss the different components of ISS.	[10]	CO1	L2

Each Components carriers with diagram 2 marks(5*2=10)

1. **Components of an Intelligent Storage System (ISS)**

- ISS is a feature-rich RAID-array that provides highly optimized I/O-processing capabilities.
- To improve the performance, storage-device provides
  - → large amount of cache
  - → multiple paths (storage-system --> storage-device)
- It handles the management, allocation, and utilization of storage-capacity.
- A storage-device consists of 4 components (Figure 1-21):
  1. Front-end
  2. Cache
  3. Back-end and
  4. Physical-disk (or simply the disk).
- A RW-request is used for reading and writing of data from the disk**.**
  1. Firstly, a read-request is placed at the host.
  2. Then, the read-request is passed to front-end, then to cache and then to back-end.
  3. Finally, the read-request is passed to disk.
- A read-request can be serviced directly from cache if the requested-data is available in cache.

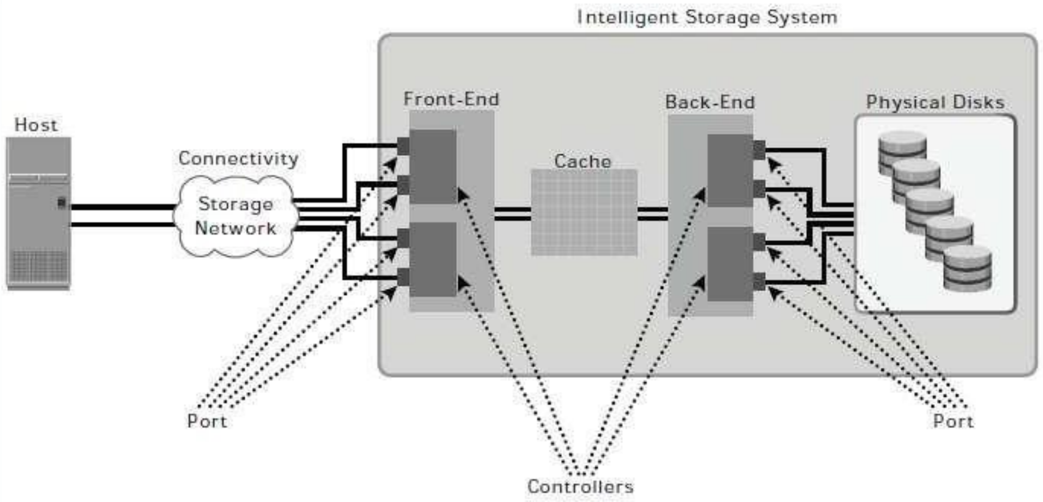

Figure 1-21: Components of an intelligent storage system

1. **Front End**

- Front-end provides the interface between host and storage.
- It consists of 2 components: 1) front-end port and 2) front-end controller.

1. Front-End Port

- Front-end port is used to connect the host to the storage.
- Each port has processing-logic that executes appropriate transport-protocol for storage-connections
- Transport-protocol includes SCSI, FC, iSCSI and FCoE.
- Extra-ports are provided to improve availability.

2. Front-End Port

- Front-end port
  → receives and processes I/O-requests from the host and
  → communicates with cache.
- When cache receives write-data, controller sends an acknowledgment back to the host.
- The controller optimizes I/O-processing by using command queuing algorithms.

1. Cache

- Cache is a semiconductor-memory.
- Advantages:
  1. Data is placed temporarily in cache to reduce time required to service I/O-requests from host For example:
  Reading data from cache takes less time when compared to reading data directly from disk (Analogy: Travelling from Chikmagalur to Hassan takes less time when compared to travelling from Dharwad to Hassan. Thus, we have
  Host = Hassan   Cache = Chikmagalur
  Disk = Dharwad).

2. Performance is improved by separating hosts from mechanical-delays associated with disks. Rotating-disks are slowest components of a storage. This is '.' of seek-time & rotational-latency

1. Back End

- The back-end

→ provides an interface between cache and disk.
→ controls data-transfer between cache and disk.

- Write operation:
  From cache, data is sent to the back-end and then forwarded to the destination-disk.
- It consists of 2 components: 1) back-end ports and 2) back-end controllers.

1. Back End Ports

- Back End Ports is used to connect the disk to the cache.

2. Back End Controllers

- Back End Controllers is used to route data to and from cache via internal data-bus.
- The controller
  → communicates with the disks when performing read- and write-operations and
  → provides small temporary data-storage.
- The controllers provides
  → error-detection and -correction (e.g. parity)
  → RAID-functionality.

Dual Controller

- To improve availability, storage-device can be configured with dual-controllers with multiple-ports In case of a port-failure, controller provides an alternative path to disks.
- Advantage:
  1. Dual-controllers also facilitate load-balancing.

Dual Port Disk

- The availability can be further improved if the disks are also dual-ported.
  In this case, each disk-port can be connected to a separate controller.

1. Physical Disk

- A disk is used to store data persistently for future-use.
- Disks are connected to the back-end using SCSI or FC.
- Modern storage-devices provide support for different type of disks with different speeds.
- Different type of disks are: FC, SATA, SAS and flash drives (pen drive).
- It also supports the use of a combination of flash, FC, or SATA.

Storage Provisioning

- It is process of assigning storage-capacity to hosts based on performance-requirements of the hosts.
- It can be implemented in two ways: 1) traditional and 2) virtual.

Traditional Storage Provisioning

al Unit (LUN)

- The available capacity of RAID-set is partitioned into volumes known as **logical-units (LUNs)**.
- The logical-units are assigned to the host based on their storage-requirements.
- For example (Figure 1-26)
  LUNs 0 and 1 are used by hosts 1 and 2 for accessing the data.

- LUNs are spread across all the disks that belong to that set.
- Each logical-unit is assigned a unique ID called a **logical-unit number (LUN#)**.
- Advantages:
  1. LUNs hide the organization and composition of the set from the hosts.
  2. The use of LUNs improves disk-utilization. For example,

Without using LUNs, a host requiring only 200 GB will be allocated an entire 1 TB disk.

With using LUNs, only the required 200 GB will be allocated to the host. This allows the remaining 800 GB to be allocated to other hosts.

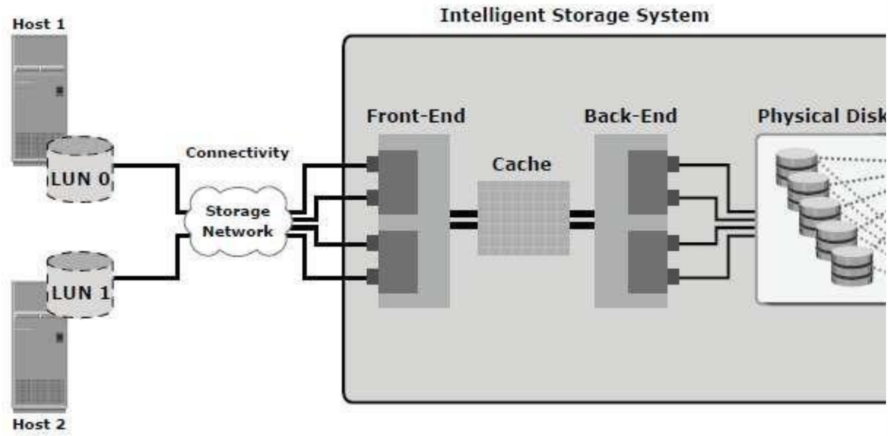

Figure 1-26: Logical-unit number

6	Explain FC connectivity options with relevant diagram.	[10]	CO2	L2

Topic: FC Connectivity

The FC architecture supports three basic interconnectivity options: point-to-point, arbitrated loop, and Fibre Channel switched fabric.

Each FC Connectivity with diagram carries 3 Marks (3*3=9)

1 marks for explanation

Point-to-Point

Point-to-point is the simplest FC configuration — two devices are connected directly to each other, as shown in Figure 5-6. This configuration provides a dedicated connection for data transmission between nodes. However, the point-to-point configuration offers limited connectivity, because only two devices can communicate with each other at a given time. Moreover, it cannot be scaled to accommodate a large number of nodes. Standard DAS uses point-to-point connectivity.

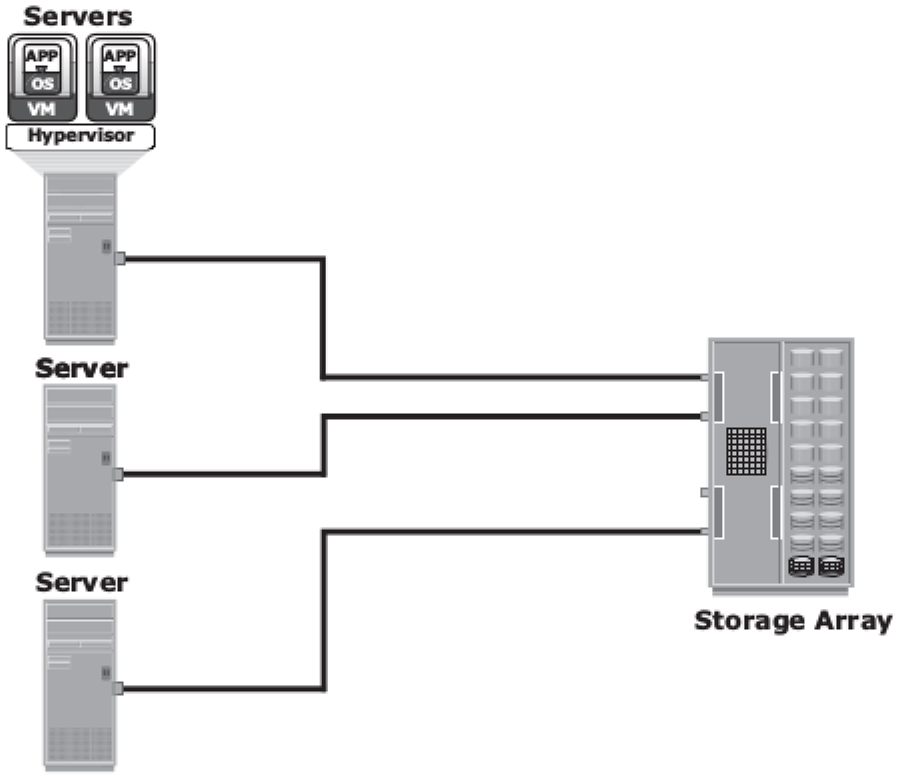

Figure 5-6: Point-to-point connectivity

Fibre Channel Arbitrated Loop

In the FC-AL configuration, devices are attached to a shared loop. FC-AL has the characteristics of a token ring topology and a physical star topology. In FC-AL, each device contends with other devices to perform I/O operations. Devices on the loop must “arbitrate” to gain control of the loop. At any given time, only one device can perform I/O operations on the loop (see Figure 5-7).

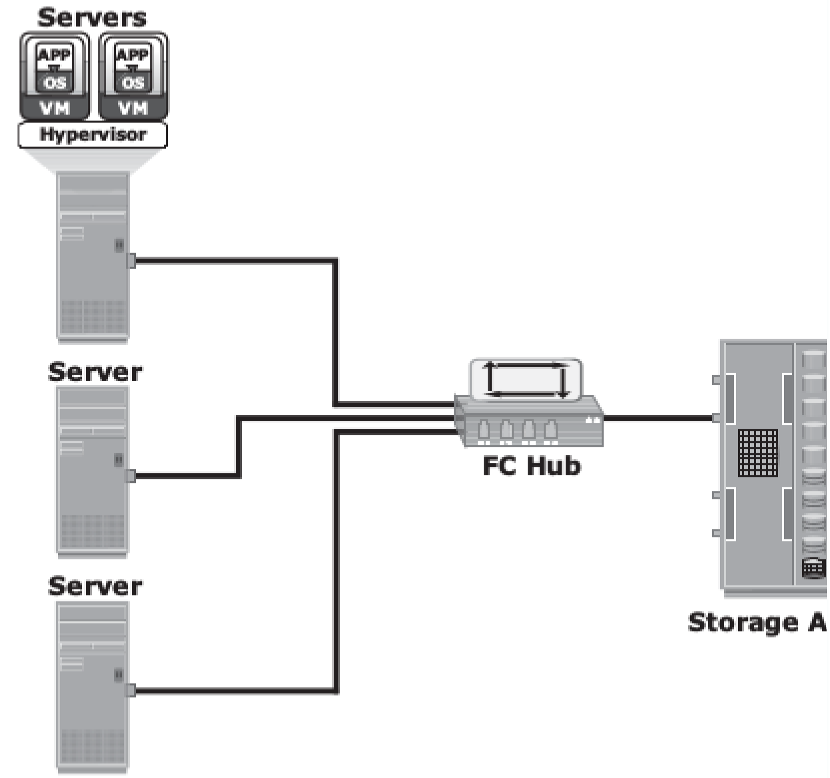

Figure 5-7: Fibre Channel Arbitrated Loop

As a loop configuration, FC-AL can be implemented without any interconnecting devices by directly connecting one device to another two devices in a ring through cables.

However, FC-AL implementations may also use hubs whereby the arbitrated loop is physically connected in a star topology.

The FC-AL configuration has the following limitations in terms of scalability:

- FC-AL shares the loop and only one device can perform I/O operations at a time. Because each device in a loop must wait for its turn to process an I/O request, the overall performance in FC-AL environments is low.
- FC-AL uses only 8-bits of 24-bit Fibre Channel addressing (the remaining 16-bits are masked) and enables the assignment of 127 valid addresses to the ports. Hence, it can support up to 127 devices on a loop. One address is reserved for optionally connecting the loop to an FC switch port. Therefore, up to 126 nodes can be connected to the loop.
- Adding or removing a device results in loop re-initialization, which can cause a momentary pause in loop traffic.

Fibre Channel Switched Fabric

Unlike a loop configuration, a Fibre Channel switched fabric (FC-SW) network provides dedicated data path and scalability. The addition or removal of a device in a switched fabric is minimally disruptive; it does not affect the ongoing traffic between other devices. FC-SW is also referred to as fabric connect. A fabric is a logical space in which all nodes communicate with one another in a network. This virtual space can be created with a switch or a network of switches. Each switch in a fabric contains a unique domain identifier, which is part of the fabric addressing scheme.

In FC-SW, nodes do not share a loop; instead, data is transferred through a dedicated path between the nodes. Each port in a fabric has a unique 24-bit Fibre Channel address for communication. Figure 5-8 shows an example of the FC-SW fabric.
In a switched fabric, the link between any two switches is called an Interswitch link (ISL). ISLs enable switches to be connected together to form a single, larger fabric. ISLs are used to transfer host-to-storage data and fabric management traffic from one switch to another. By using ISLs, a switched fabric can be expanded to connect a large number of nodes.

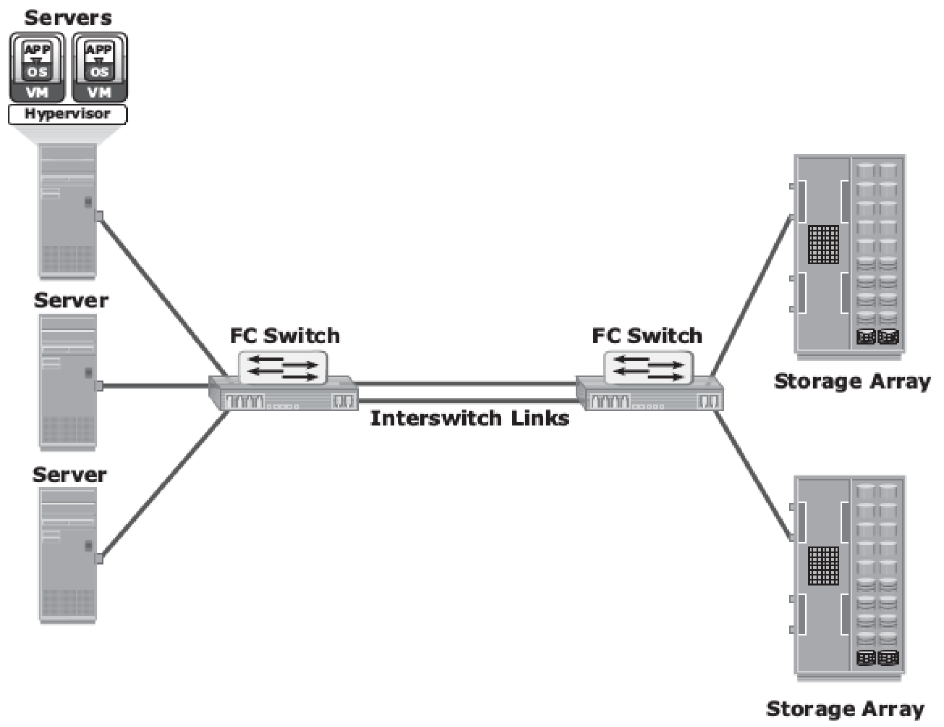

Figure 5-8: Fibre Channel switched fabric

A fabric can be described by the number of tiers it contains. The number of tiers in a fabric is based on the number of switches traversed between two points that are farthest from each other. This number is based on the infrastructure constructed by the fabric instead of how the storage and server are connected across the switches.

When the number of tiers in a fabric increases, the distance that the fabric management traffic must travel to reach each switch also increases. This increase in the distance also increases the time taken to propagate and complete a fabric reconfiguration event, such as the addition of a new switch or a zone set propagation event. Figure 5-9 illustrates two-tier and three-tier fabric architecture.

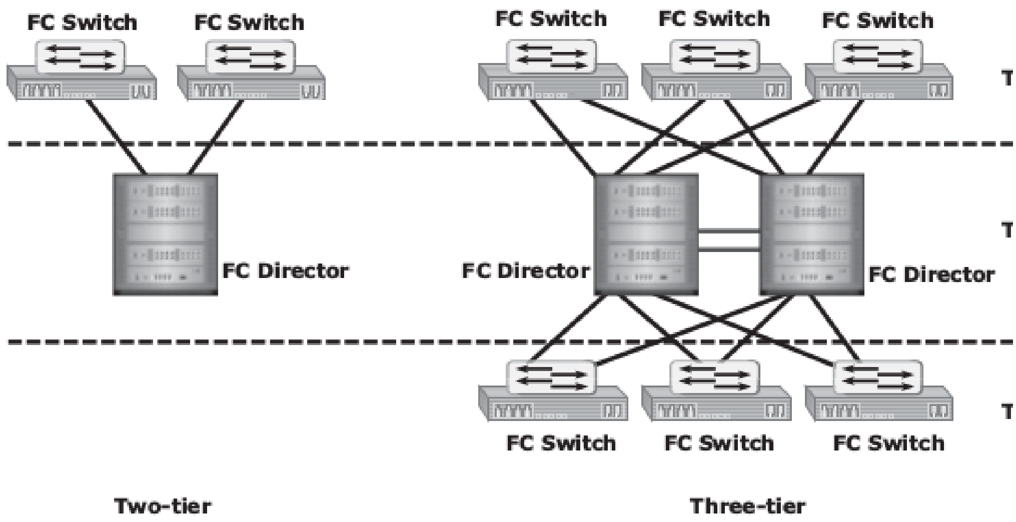

Figure 5-9: Tiered structure of Fibre Channel switched fabric

FC-SW Transmission

FC-SW uses switches that can switch data traffic between nodes directly through switch ports. Frames are routed between source and destination by the fabric. As shown in Figure 5-10, if node B wants to communicate with the node D, the nodes should individually login first and then transmit data via the FC-SW. This link is considered a dedicated connection between the initiator and the target.

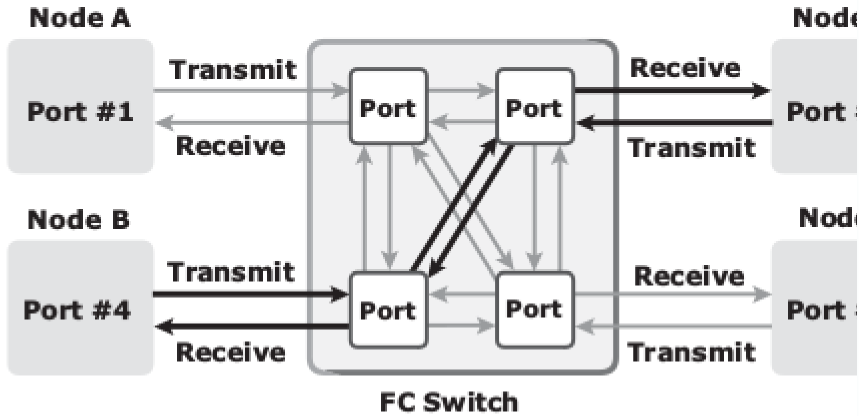

Figure 5-10: Data transmission in Fibre Channel switched fabric