Internal Assessment Test 1 – Sept. 2018

| Sub: | Storage Area Networks | | | | Sub Code: | 15CS754 | Branch: | ISE / CSE |
|---|---|---|---|---|---|---|---|---|
| Date: | 20-09-18 | Duration: | 90 min's | Max Marks: 50 | Sem / Sec: | VII SEM ISE A & B / CSE A ,B & C | | OBE |

## SCHEME OF EVALUATION

| | | MARKS | CO | RBT |
|---|---|---|---|---|
| 1. | Explain the architecture and evolution of storage technology with a neat diagram. | [10] | CO1 | L1 |

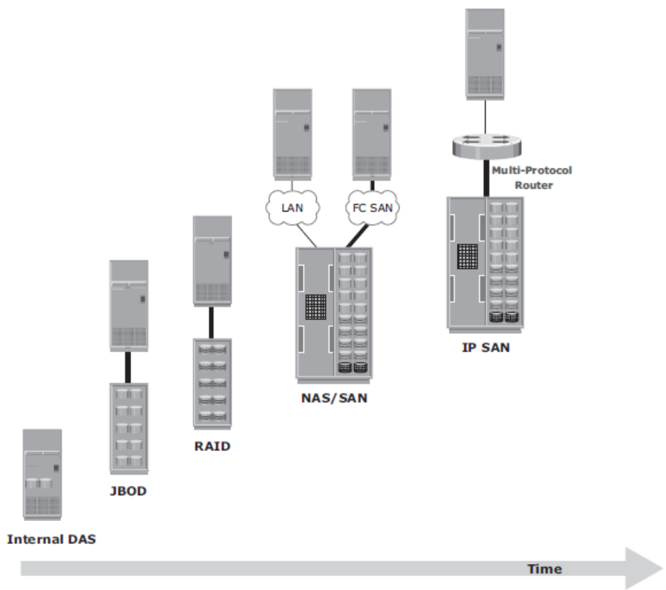

Evolution of storage architectures

**1. Redundant Array of Independent Disks (RAID):** This technology was developed to address the cost, performance, and availability requirements of data.
**2. Direct-attached storage (DAS):** This type of storage connects directly to a server (host) or a group of servers in a cluster. Storage can be either internal or external to the server..
4MM
**3. Storage area network (SAN):** This is a dedicated, high-performance Fibre Channel (FC) network to facilitate block-level communication between servers and storage.
**4. Network-attached storage (NAS):** This is dedicated storage for file serving applications.
**5. Internet Protocol SAN (IP-SAN):** One of the latest evolutions in storage architecture, IP-SAN
is a convergence of technologies used in SAN and NAS. IP-SAN provides block-level communication across a local or wide area network (LAN or WAN), resulting in greater consolidation and availability of data

2. Discuss the key characteristics of a data center, with a neat diagram. [10] CO1 L2


- Availability
- Security
- Scalability
- Performance
- Data Integrity
- Capacity
- Manageability

3. (a) Explain disk drive components with suitable diagram. [5] CO1 L1

- ***<u>Platter:</u>*** A typical HDD consists of one or more flat circular disks called *platters*. The data is recorded on these platters in binary codes (0s and 1s). The set of rotating platters is sealed in a case, called *Head Disk Assembly* (HAD)
- ***<u>Spindle:</u>*** A spindle connects all the platters and is connected to a motor. The motor of the spindle rotates with a constant speed
- ***<u>Read/Write Head:</u>*** *Read/Write (R/W) heads,* read and write data from or to a platter. Drives have two R/W heads per platter, one for each surface of the platter. The R/W head changes the magnetic polarization on the surface of the platter w
- ***<u>Actuator Arm Assembly:</u>*** Each platter has two R/W heads, one for each surface. R/W heads are mounted on the *actuator arm assembly* , which positions the R/W head at the location on the platter where the data needs to be written
- ***<u>Drive Controller:</u>*** The controller is a printed circuit board, mounted at the bottom of a disk drive. It consists of a microprocessor, internal memory, circuitry, and firmware. The firmware controls the power to the spindle motor and the speed of the motor

(b) Consider a Disk I/O System in which an I/O request arrives at the rate of 80 IOPS. Service time is 6 ms. Compute the following [5] CO1 L3

(i) Utilization (ii) Total Response Time.

(iii) Average Queue Size (iv) Time Spent by a request in a queue

Arrival Rate= a=80 IOPS
Average service Time=Rs= 6ms
Hence, the utilization of I/O controller (U) is given by:
U = a * Rs
= 80 * 6 * 10^-3
= 0.48

**Total Response Time (R) is given by:**

$R = \frac{R_s}{1-U}$

$= \frac{6}{1-0.48}$

= 11.538 ms

Average queue size $= \frac{U*U}{(1-U)}$

$= \frac{0.48*0.48}{(1-0.48)}$

= 0.44

Thus, the total time spent by a request in a queue
= U * R
=0.48 * 11.538
=5.538 ms

Now if the controller power is doubled, the service time is halved. Consequently, Rs= 3ms in this scenario.
Therefore, a=80 IOPS
Utilization (U) = a * Rs
= 80 * 3 * 10^-3
= 0.24

Total Response Time (R) $= \frac{R_s}{1-U}$

$= \frac{3}{1-0.24}$

= 4 ms

Average Queue Size $= \frac{U*U}{(1-U)}$

$= \frac{0.24*0.24}{1-0.24}$

= 0.0757

**Hence, Time spent by request in a queue = U * R= 0.24 * 4= 0.96ms.**

4. Explain the following with appropriate diagrams [10] CO1 L1

(i) RAID 0 (ii) RAID 1 (iii) Nested RAID (iv) Comparison of RAID Levels

**RAID 0**

- **Data is distributed across the HDDs in the RAID set.**
- **Allows multiple data to be read or written simultaneously, and therefore improves performance.**
- **Does not provide data protection and availability in the event of disk failures.**

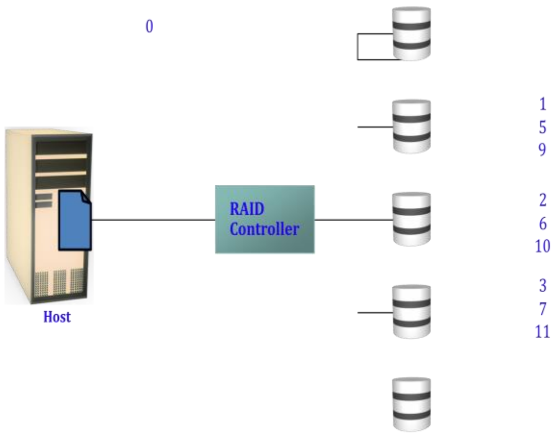


**RAID 1**

- **Data is stored on two different HDDs, yielding two copies of the same data.**
  - **Provides availability.**
- **In the event of HDD failure, access to data is still available from the surviving HDD.**
- **When the failed disk is replaced with a new one, data is automatically copied from the surviving disk to the new disk.**
  - **Done automatically by RAID the controller.**
- **Disadvantage: The amount of storage capacity is twice the amount of data stored.**
- **Mirroring is NOT the same as doing backup!**

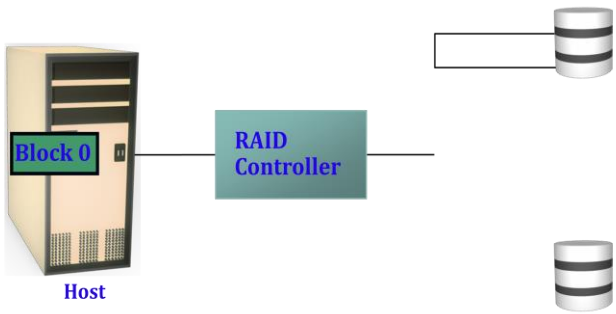


## NESTED RAID

- Combines the performance benefits of RAID 0 with the redundancy benefit of RAID 1.
- RAID 0+1 – Mirrored Stripe
  - Data is striped across HDDs, then the entire stripe is mirrored.
  - If one drive fails, the entire stripe is faulted.
  - Rebuild operation requires data to be copied from each disk in the healthy stripe, causing increased load on the surviving disks.
- RAID 1+0 – Striped Mirror
  - Data is first mirrored, and then both copies are striped across multiple HDDs.
  - When a drive fails, data is still accessible from its mirror.
  - Rebuild operation only requires data to be copied from the surviving disk into the replacement disk.

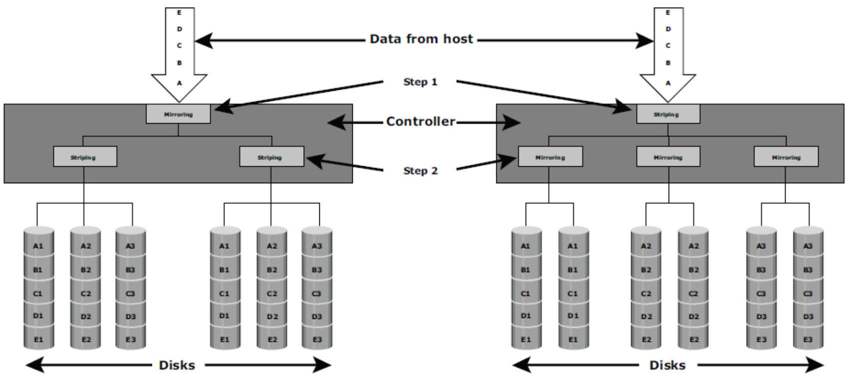


(a) RAID 1+0 (b) RAID 0+1

## COMPARISION OF RAID LEVELS

| RAID | Min Disks | Storage Efficiency % | Cost | Read Performance | Write Performance |
|---|---|---|---|---|---|
| 0 | 2 | 100 | Low | Very good for both random and sequential read | Very good |
| 1 | 2 | 50 | High | Good<br>Better than a single disk | Good<br>Slower than a single disk, as every write must be committed to two disks |
| 3 | 3 | (n-1)*100/n where n= number of disks | Moderate | Good for random reads and very good for sequential reads | Poor to fair for small random writes<br>Good for large, sequential writes |
| 5 | 3 | (n-1)*100/n where n= number of disks | Moderate | Very good for random reads<br>Good for sequential reads | Fair for random write<br>Slower due to parity overhead<br>Fair to good for sequential writes |
| 6 | 4 | (n-2)*100/n where n= number of disks | Moderate but more than RAID 5 | Very good for random reads<br>Good for sequential reads | Good for small, random writes<br>(has write penalty) |
| 1+0 and 0+1 | 4 | 50 | High | Very good | Good |

5. Explain the features of High End Storage Systems, with a diagram [10] CO1 L4

- Also referred as Active-active arrays
  - I/O’s are serviced through all the available path
- Following are high-end array capabilities:
  - Large storage capacity
  - Huge cache to service host I/Os
  - Fault tolerance architecture
  - Multiple front-end ports and support to interface protocols
  - High scalability
  - Ability to handle large amounts of concurrent I/Os
- Designed for large enterprises

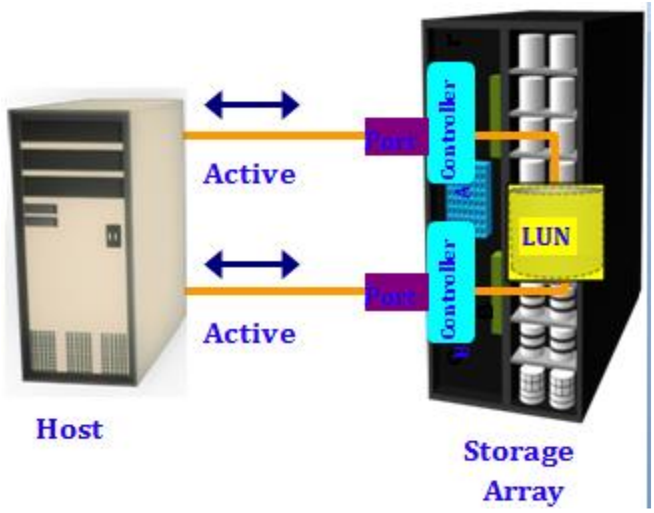


6. With a neat diagram , Explain the structure of read and write operations in cache [10] CO1 L1

READ OPERATION FROM CACHE

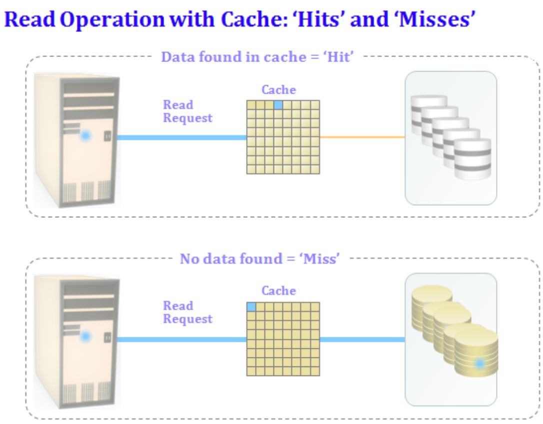

WRITE OPERATION TO CACHE

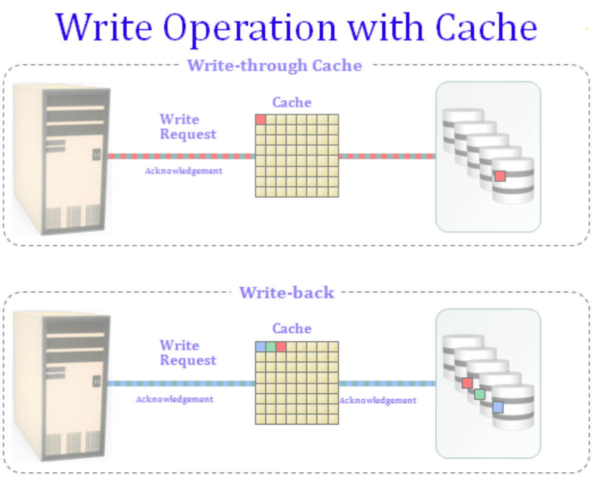


| | | | | |
|---|---|---|---|---|
| 7. | Explain the connectivity options of FC architecture with a neat diagram | [10] | CO1 | L1 |

The FC architecture supports three basic interconnectivity options: point-to-point, arbitrated loop (FC-AL), and fabric connect.

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

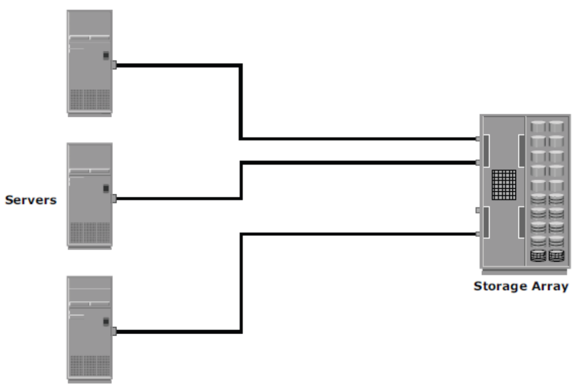


Point-to-point topology

In the FC-AL configuration, devices are attached to a shared loop, as shown in Figure below . 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. As a loop configuration, FC-AL can be implemented without any interconnecting devices by directly connecting one device to another in a ring through cables

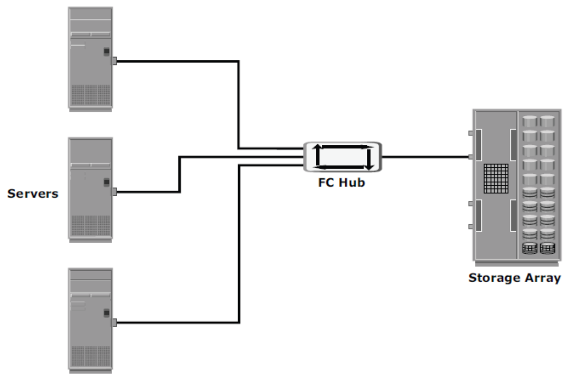

Fibre Channel arbitrated loop

1. FC-AL shares the bandwidth in the loop. Only one device can perform I/O operations at a time. Because each device in a loop has to wait for its turn to process I/O request, the speed of data transmission is low in an FC-AL topology.
2. FC-AL uses 8-bit addressing. It can support up to 127 devices on a loop.
3. Adding or removing a device results in loop re-initialization, which can cause a momentary pause in loop traffic.

8. With a neat diagram, Explain the components of Intelligent Storage Systems (ISS) [10] | CO1 | L1

Intelligent Storage Systems are RAID arrays that are:

- Highly optimized for I/O processing
- Have large amounts of cache for improving I/O performance
- Have operating environments that provide:
  – Intelligence for managing cache
  – Array resource allocation
  – Connectivity for heterogeneous hosts
  – Advanced array based local and remote replication options

Intelligent storage system provides several benefits over a collection of disks in an array (JBOD) or even a RAID arrays:

  – Increased capacity
  – Improved performance
  – Easier data management
  – Improved data availability and protection
  – Enhanced Business Continuity support
  – Improved security and access control

Components of ISS