Internal Assessment Test 1 – September 2019

Sub:	**Storage Area Networks**					Sub Code:	**15CS754**	Branch:	**ISE**
Date:	21-09-19	Duration:	90 min's	Max Marks:	50	Sem / Sec:	VII SEM ISE A		OBE

SCHEMES AND SOLUTIONS

1. **Explain the architecture and evolution of storage technology with a neat diagram.**

Data created by individuals/businesses must be stored for further processing
Type of storage used is based on the type of data and the rate at which it is created and used
Examples:

a. Individuals: Digital camera, Cell phone, DVD's, Hard disk
b. Businesses: Hard disk, external disk arrays, tape library

Storage model: An evolution

c. Centralized: mainframe computers
d. Decentralized: Client –server model
e. Centralized: Storage Networking

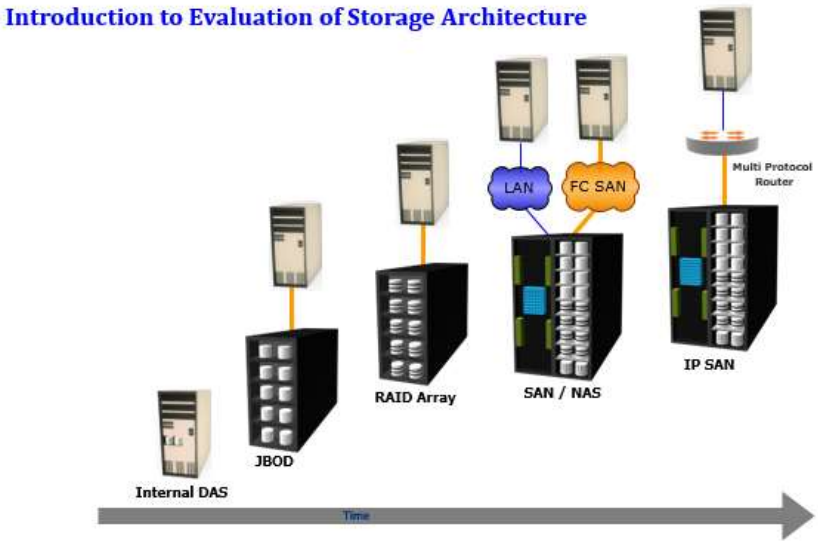

Redundant Array of Independent Disks (RAID): This technology was developed to address the cost, performance, and availability requirements of data. It continues to evolve today and is used in all storage architectures such as DAS, SAN, and so on.
Direct-attached storage (DAS): This type of storage connects directly to a server (host) or a group of servers in a cluster. Storage can be either internal or external to the server. External DAS alleviated the challenges of limited internal storage capacity.
Storage area network (SAN): This is a dedicated, high-performance *Fibre Channel (FC) network to facilitate block-level communication between servers and storage. Storage is partitioned and assigned to a server for accessing its data. SAN offers scalability, availability, performance, and cost benefits compared to DAS.*
Network-attached storage (NAS): *This is dedicated storage for file serving applications. Unlike a SAN, it connects to an existing communication network (LAN) and provides file access to heterogeneous clients. Because it is purposely built for providing storage to file server applications, it offers higher scalability, availability, performance, and cost benefits compared to general purpose file servers.*
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.**

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

3**. A. Explain disk drive components with suitable diagram.**

Platter: 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)

Spindle: A spindle connects all the platters and is connected to a motor. The motor of the spindle rotates with a constant speed

Read/Write Head: *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

Actuator Arm Assembly: 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

Drive Controller: 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

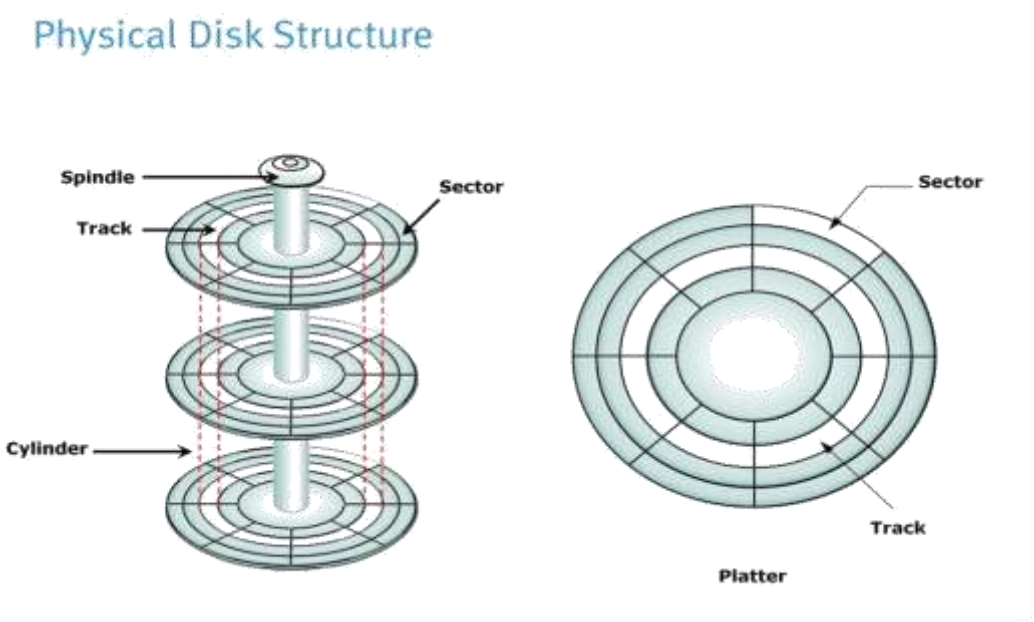

3. B**. Consider a Disk I/O System in which an I/O request arrives at the rate of 80 IOPS. Service time is 6 ms. Compute the following**

(i) Utilization (ii) Total Response Time.

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

Arrival rate (a) = 80 IOPS,
consequently, the arrival time Ra = 1/a = 1/80 = 12.5 ms
Rs = 6 ms (given)
Utilization (U) = Rs/Ra = 6/12.5 = 0.48 or 48%
Response time (R) = Rs/(1 -U) = 6/(1-0.48)= 11.5 ms
Average queue size = U /(1-U) = (0.48) 2 /(l-0.48) = 0.44
Time spent by a request in a queue = UxR,
or the total response time-service time = 0.48x1 1.5 = 5. 52 ms

1 . Utilization (U) = Rs/Ra = 3/12.5 = 0.24 or 24 %
2. Response time (R) = Rs/(1-U) = 3/(1-0.24) = 3.9 ms
3. Average queue size = U 2 /(l-U) = (0.24) 2 /(1-0.24) = 0.08
4. Time spent by a request in a queue = 0.24x3.9 = 0.936 ms

4. Discuss the features of High End Storage Systems, with a diagram

- 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

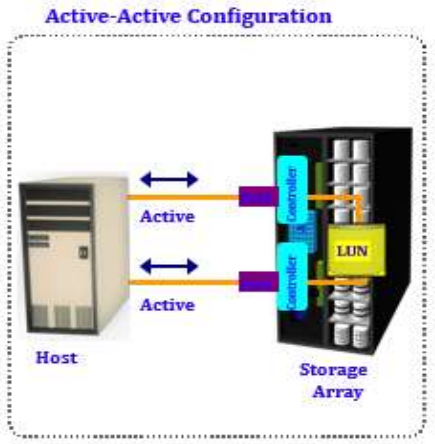

5. **With a neat diagram , Explain the structure of read and write operations in cache**

Write Operation with Cache

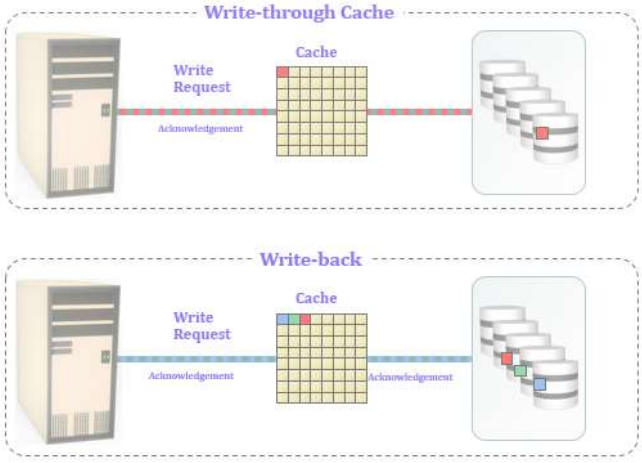

Read Operation with Cache: 'Hits' and 'Misses'

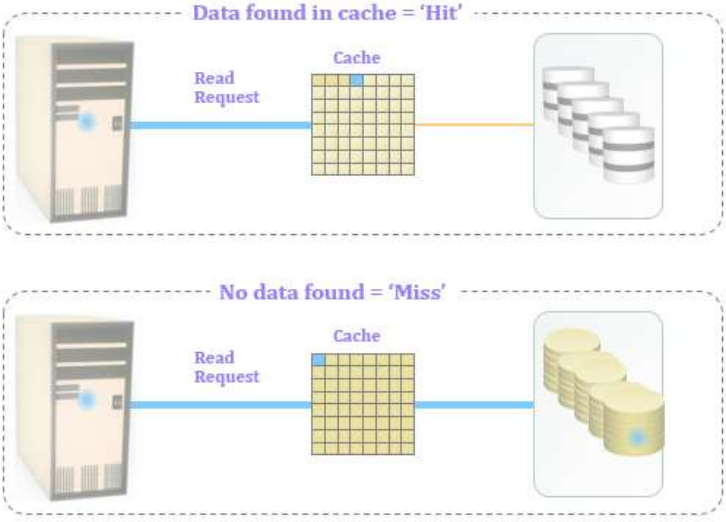

6. **Explain the following with appropriate diagrams**

(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.

1. Striping

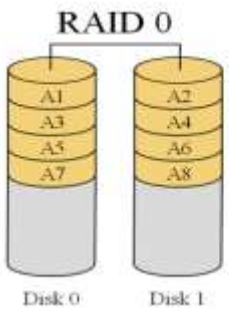

2. Mirroring

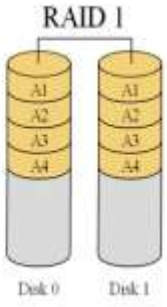

COMPARISON OF RAID LEVELS

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 Better than a single disk	Good 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 Good for large, sequential writes
5	3	(n-1)*100/n where n= number of disks	Moderate	Very good for random reads Good for sequential reads	Fair for random write Slower due to parity overhead 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 Good for sequential reads	Good for small, random writes (has write penalty)
1+0 and 0+1	4	50	High	Very good	Good

7. **With a neat diagram, Explain the components of Intelligent Storage Systems (ISS)**

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