

EE4217 Quiz 2 (Total marks 35%)

1. Hardware RAID is implemented by intelligent storage systems external to the host, or, at minimum, intelligent controllers in the host that offload the RAID management functions from the host.
2. Software RAID usually describes RAID that is managed by the host. Typically it is implemented via Logical Volume Manager on the host.
3. Access to, and manipulation of data on tape can be a time consuming process. As data is stored sequentially on tape, random access is not possible.
4. Technologies and strategies were developed so that data can be stored centrally and allowed access to it for different applications. There are two main ways to access data in Networked Storage – SAN (Storage Area Network) and NAS (Network Attached Storage).
5. With the emergence of Optical Disks came the advantage of “write-protected” data and random access. Large quantities of these disks were assembled into optical jukeboxes.
6. A SCSI device that starts a communication is an initiator. A SCSI device that services a request is a target. The Logical Unit Numbers (LUNs) identifies a specific Logical Unit in a target.
7. When looking at particular devices, SCSI can support multiple devices simultaneously while IDE/ATA can only support a single device at a time.
8. Fibre Channel is a high-speed interconnect used in networked storage to connect servers to shared storage devices. Fibre Channel components include HBAs, hubs, switches, cabling, and disks.
9. Disk platters spin at speeds of several thousand revolutions per minute and multiple platters are connected by a spindle.
10. When the spindle rotation has stopped, the air gap is removed and the R/W heads rest on the surface of the platter in a special area near the spindle called a landing zone.
11. Disk data is recorded in tracks, and it is divided into sectors which typically hold 512 bytes of user data.
12. Partitioning divides the disk into logical containers (known as volumes), each of which can be used for a particular purpose.

13. Mean Time Between Failure (MTBF) is the amount of time that one can anticipate a device to work before an incapacitating malfunction occurs.
14. RAID 0+1 is one way of combining the speed of RAID 0 with the redundancy of RAID 1. RAID Level 3 stripes data for high performance and uses parity for improved fault tolerance.
15. Like hot spares, hot swaps enable a system to recover quickly in the event of a failure. With this setup, the user can replace the failed hardware (such as a controller) without having to shut down the system.
16. The read cache hit ratio (or hit rate), usually expressed as a percentage, describes how well the read cache is performing and it is obtained by dividing the number of read cache hits by the total number of read requests.
17. The EMC CLARiiON architecture includes fully redundant, hot swappable components—meaning the system can survive the loss of a fan or a power supply, and the failed component can be replaced without powering down the system.
18. The EMC Symmetrix DMX series arrays deliver the highest levels of performance and throughput for high-end storage. It incorporates the following features: 1) Direct Matrix Interconnect 2) Dynamic Global Memory 3) Engenuity Operating Environment 4) Advanced processing power 5) High-performance back end 6) A fully fault-tolerant design.
19. Zoning is a switch function that allows nodes within the fabric to be logically segmented into groups that can communicate with each other. Its function controls the process at login by only letting ports in the same zone to establish link level services.
20. A more flexible method called WWN zoning uses the unique identifiers of a node which have been recorded in the switches to either allow or block access.