(1)With the help of suitable schematic diagrams, show how the processing, storage and networking functions in earlier generation storage systems used to be different than present day storage systems.

Consider the example of how a remote user would access data stored on a machine via an application program in the two systems to illustrate how the processing, storage and networking functions have gotten separated.

(2)Compare the pros and cons of accessing data stored on the nodes of a computer network with NFS APIs and with Cloud Storage APIs.

(3)Consider a real-time storage system, i.e a storage system in which user requests have to satisfy real-time constraints. Assume that the latencies of Hard Disk, Main Memory and Network link are d, m, n sec per kB of data resp. and that there are n users whose data request size and maximum waiting time are r(i) kB and T(i) sec resp. Give a simple algorithm for ordering the requests so as to satisfy real-time constraints. State any assumptions made in your algorithm.

(4)Why is a Dropbox type of storage system different from a desktop filesystem?

(5)Why is a storage system for holding mail different from a file system for holding crawled pages of a search engine?

(6)Suppose I have a tree-based information system. If I do not have update-in-place property, what is the worst case time for updating some node in the system? Suppose I want to do better than this worst case. What can I do?

(7)Explain why there are 2 FATs in early floppy-based designs. If you instead now have a SSD, comment on the suitability of such designs and what you would change.

(8)If I want to send 8KB write blocks over NFS on wireless, and the BER is 10E-6, how many expected retries are needed if the retry is at the appl level (8K write blocks) vs if retries are at the netw packet

level with pkt size of 1KB?

(9)An interrupt handler cannot block. Argue carefully why this is an important constraint. How can you satisfy this in practice? In a multiprocessor OS, identify at least 2 types of blockage and how workarounds for this.

(10)If you hold a spinlock, you should not sleep. Why?

(11)What is the difference between a process context, a system context and an interrupt context. Give an example of each.

(12)Why is the "old" blocklayer not so good for flash devices? Which specific part would you modify/throw out?

(13)Why is the longevity of digital data a difficult problem? Give a plausible solution to each of the subproblems involved.

(14)Write a simple device driver that can be used to read and manage the values logged by the kernel for some specific purpose.

(15)What is the difference b/w a Networked FileSystem, Cluster Filesystem and Distributed Filesystem? Give the names of two popular filesystems that you know for each type. Give 3 example cases to explain why implementing POSIX semantics over NFS or Cluster or Distributed filesystem is a difficult problem.

(16)In a Networked Filesystem, usually replicas of some data object are maintained on various Nodes to improve availability and performance. Give an example of a scenario in which inconsistency could arise among the replicas. How does NFSv2 handle inconsistency among replicas? NFSv3? What is the difference between the NFSv3 approach and the NFSv4 approach?

(17)NFSv2 uses deamons like Portmapper, mountd, lockd, statd. Suppose a modern Storage Area Network which has 5 subnets seperated by firewalls runs on the NFSv2 filesytem protocol. Suppose a user in Subnet-1 wishes

to update a file located on Subnet-4. Explain how the NFSv2 deamons would function in order to offer security for the transaction. What kind of simple security gaurantees could be gauranteed for the transaction? Give a few cases where security would break down.

(18)"NFSv2 is a stateless protocol while NFSv4 is a stateful protocol". Explain this statement with the help of typical protocol operation and data structure.

(19)Fibre Channel/10GB Eth/TCP Q1 : Describe the "buffer credits" based Traffic management and Flow Control mechanisms of the Fibre Channel protocol. Explain why FC protocol is lossless. Is FC an effective protocol for a multi-Data Center SAN which has Switches, Routers, Multiplexsers in transmission path that introduce delays? Is FC a good choice for a WAN?

(20)Suppose a computer program has generated a new 4KB page of data during computation that may need to be accessed again. It can either be kept in DRAM memory or after some time "k", written out to less expensive memory. Find break-even time "k" for a multi-tier Storage system which has 5% SSD, 10% DRAM, 25% HDD, 60% Tape if their cost/GB are resp. Rs.2500, Rs.750, Rs.200, Rs.38 and cost/iops are resp. Rs.1, Rs.5, Rs.50 and Rs.300. In practice, new data generation and memory I/O activity during a computation varies over time. Suggest what could be done to achieve the best cost-performance values at any given time, in such a case?

(21)Explain why Consistency, Availability and Performance tend to be contradictory forces in a distributed filesystem. Taking the example of a system that uses a Log-structured distributed filesystem, explain how it can be designed so as to ensure all the three - Consistency, Availability and Performance.

(22)In a Google File System, 5 GFS clients simultaneously access the same file located on the nodes of GFS. Assume that two of the requests are Write requests while the others are Read requests. If large chunk size like 64MB is used, will the chunks become hotspots? What can be done to solve this problem? What

can be done to avoid locking/deadlock problems, if any? Is there a possibility that stale data is served and how can this be fixed? Give all your answers with the help of Architecture diagrams showing control messages like Read/Write/Metadata-Request etc.

(23)What is meant by "Zero-copy" semantics w.r.t the operation of FC protocol in a Storage Area Network? Taking the example of a NFS Client reading a file, show how "page flipping" zero-copy technique is more advantageous than 2-copy.

(24)Explain what provision is made for data packet loss or misordering during transit in each of the following SAN Transport layer protocols - FC, FC over IP, 10G Ethernet, iSCSI, Infiniband, RDMA? Hence discuss the suitability of each protocol for the cases of high data throughput, QoS, multiple-node Data Center SAN interconnected by Switches & Routers, better Fabric/Network link BW utilization. State any assumptions that you may need to make.

(25)Deduplication is a popular feature in disk-based backup systems. Suppose that Sunil has 10 copies of a 5MB spreadsheet file on his linux v.2.6 computer system, some of which are identical files and some are very similar files with just a few words changed. Give the pseudo-code of a file-level deduplication algorithm. Assume any data structures like a dedupe index file, pointers etc that you may need. Next, give the pseudo-code of a block-level deduplication algorithm.
What is the savings? If Sunil moves to a NFS system, how will the pseudocode algorithms change?