

---

**Name:** \_\_\_\_\_

**Instructions**

There are six (6) questions on the exam. You may find questions that could have several answers and require an explanation or a justification. As we've said, many answers in storage systems are "It depends!". In these cases, we are more interested in your justification, so make sure you're clear. Good luck!

If you have several calculations leading to a single answer, please place a 

box around your answer
------------------------

.

**Problem 1 : Short answer. [30 points]**

- (a) CD and DVD media have a continuous spiral track whereas hard drive media are organized into thousands of concentric tracks. Given a 4-head, 15,000 RPM hard drive with a 1.0 ms track switch time and 1.5 ms cylinder switch time, calculate the percentage improvement in the sustained media transfer rate between (i) the above design and (ii) an alternate design with a single head and a continuous spiral track.
  
- (b) Disk-based video recording and playback systems use large I/O requests (e.g., 1 MB and up). Why?
  
- (c) When is DMA a bad choice for communicating with an I/O controller (as compared to programmed I/O)?
  
- (d) Your friend is very excited about their new disk, because it has a one million hour MTBF rating. Explain why your friend can't expect to never have to buy another disk drive, even if they don't run out of capacity.

- (e) Many file systems implement directories as unordered lists of entries. Explain why this is a problem when a directory has many entries.
- (f) Modern disks continue reading sectors into their on-board memory after fetching those requested by the host. Why does this improve performance in many systems?
- (g) Why do the buffers in many disk drives use a single, dedicated segment for writes?
- (h) You have been asked to buy a new desktop computer to be used as the main Carnegie Mellon web server. So far, you've selected a system with a 1 GHz processor and a 3.5-inch, 10,000 RPM disk. You have a little extra money that you can spend on one of three things: (i) a faster processor, (ii) a 2.5-inch disk (with the same capacity), or (iii) a 15,000 RPM disk (with the same capacity). Which do you upgrade? Justify your answer.
- (i) What does *fairness* mean in the context of disk drive scheduling algorithms?
- (j) Is it possible to recalculate the information in a file system superblock if the block were corrupted? How, or why not?





Local area networks (LANs) avoid the connection setup delay by transferring data immediately without first establishing a connection. This is known as a *packet-switched* service. Because network resources are not reserved for any individual connection, it is possible for data to be lost en route from the sender to the receiver (e.g., because of congestion). When this happens, the data must be retransmitted by the sender after a timeout.

The maximum TCP/IP packet size using Ethernet is 1500 bytes: 0–1460 bytes for user data and 40 bytes of overhead (20 bytes for the TCP headers and 20 bytes for the IP headers). Assume the underlying network is a 1 Gbit/s switched Ethernet. Assume also that the retransmission timeout is five times the round-trip time between the sender to the receiver, and that any packets following a lost packet must be retransmitted as well.

(c) What is the total transmission time (including setup) for sending the following over a 3-km TCP/IP link: (i) one 512-byte sector, (ii) 1 MB?

(d) When sending 1 MB of data, what is the maximum error rate on the TCP/IP-based network (as a percentage of packets sent) before the use of Fibre Channel would be preferred?



(c) Joe heard that sometimes disk drives grow media defects and thereby destroy one or more sectors of existing disk content. He thinks that he doesn't have to worry, because he uses write-ahead logging. Is he correct? If so, explain why. If not, explain why not and propose an additional crash recovery step for addressing the problem.

(d) After recovering from one particular crash, Joe discovered a file that contained random data from a previously deleted file. Explain how this could happen in a write-ahead logging system, and suggest a mechanism Joe can use to prevent this integrity problem from happening again.





*This page intentionally left blank in case you need scratch space.*

*This page intentionally left blank in case you need scratch space.*

*This page intentionally left blank in case you need scratch space.*