

---

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

### Instructions

There are four (4) 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. [56 points]

- (a) Imagine a disk with 500 KB per track and a rotation speed of 6000 RPM. Once positioning is complete, would you expect the disk to be able to read approximately 5 MB off the media in 100ms? ("Yes" or "No") Explain why or why not.
  
- (b) Modern disk drives have a vendor specified mean time between failures (MTBF) of over 1,000,000 hours. But, most disks are expected to be used for less than 45,000 hours (approximately 5 years). Given these facts, how is a disk whose MTBF is 2,000,000 hours better than one whose MTBF is 1,000,000 hours?
  
- (c) Dr. Smith saves every email he receives by appending it to a large file. Over the last year, the file has grown to 1 GB in size. In running a program to scan the file, from beginning to end, he notices that the read bandwidth is less than 5 MB/s. But, his system's specifications indicate that the disk can provide 50 MB/s. Give one likely explanation for the low observed bandwidth in reading the file.

- (d) Some disks' firmware discards data from the on-board buffer/cache after it has been sent to the host that requested it, rather than retaining it in the cache. Explain how the firmware engineers might justify this approach.
- (e) Most systems allow a program to create a "symbolic link" (a.k.a. a "soft link") to a file, even if that file does not currently exist. Why doesn't the target file have to exist in order to create the symbolic link?
- (f) Some modern disks perform write-back caching, wherein write requests are reported complete once the corresponding data is transferred from the host into the on-board RAM. Dr. Jones understands that this can result in the disk firmware's disk scheduler reordering the media writes in order to improve efficiency, which can reverse with the file system's ordering of metadata updates. But, he proposes that changing the file system to use write-ahead logging to protect its metadata integrity solves the problem. Is he correct? ("Yes" or "No") Explain (briefly).
- (g) An engineer designing a new mobile phone is arguing that it should be equipped with a small disk drive, instead of a Flash-based SSD, since random access performance is not a concern for the expected workload. Give an argument in favor of the SSD, even in this environment.

**Problem 2 : Finding service time. [20 points]**

Find the service time for the following request stream to a very simple disk that has 10000 cylinders, 8 surfaces, 200 sectors per track, and rotates at 10000 RPM. In this case, LBNs are mapped directly to PBNs, requests are serviced in FIFO order, and seek time is a linear function of cylinder distance.

$$\text{seek\_time(ms)} = 0.0006 \times \text{cylinder\_distance} + 3$$

(Remember that a linear seek curve is not realistic, but is a simplification for this problem.) Assume that the disk starts at LBN 0 and that head switches are instantaneous. The layout mapping is shown in Figure 1.

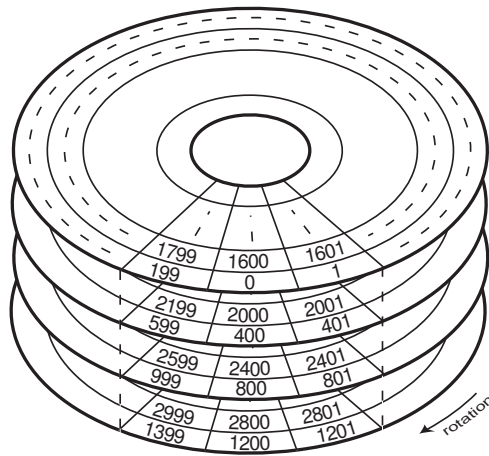


Figure 1: LBN mappings onto physical sectors on the disk.

The request stream for which you should compute the service times is shown in Table 1.

Request Stream	
LBN	Size (sectors)
0	10
4800600	20
12800200	10
6400400	20

Table 1: Request stream.

- (a) Compute the cylinder, surface and offset for the first LBN of each request in the request stream (i.e. for LBN 0, 4800600, etc.)

- (b) Calculate the total service time for the request stream by filling in the table 2. Time (ms) shows the time at which the head is at the given LBN. The LBN is mapped to a cylinder, surface, and offset within the track. The “next event” field shows which part of servicing the request (seek, rotation, or transfer) takes place next.

time (ms)	LBN	cylinder	surface	offset	next event
0	0	0	0	0	transfer 10 sectors
	10	0	0	10	seek

Table 2: *Timetable of events.*

**Problem 3 : More short answer. [24 points]**

- (a) A file system engineer implementing a log-structured file system noticed that using 128 KB segments instead of 64 KB segments nearly doubles file system write throughput. Explain why.
- (b) Imagine a new file system designed for a hybrid storage system that consists of one disk and one Flash SSD. Imagine that the workload consists of two types of files: (1) files that are randomly read in small units with no locality, and (2) large files that are read and written sequentially. Which files would you suggest assigning to each device? Explain.
- (c) Imagine an inode structure that uses 12 direct block pointers, 1 indirect block pointer, and 2 double indirect block pointers. With a 4 KB block size and 64-bit unsigned integers for block pointers, what is the largest file size?

**Problem 4 : Instructor trivia. [up to 2 bonus points]**

(a) Which instructor has kids?

(b) Which instructor is Canadian?

(c) Which instructor will take a new job in May?

(d) Which instructor most needs an obnoxious stopwatch alarm to use during their lectures?

(e) Where (city, not company) should Zoheb take a job after completing his INI M.S. program?