

Name: _____

Instructions

There are three (3) 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. [48 points]

- (a) Imagine a file system that uses write-ahead logging (WAL) to protect the integrity of its metadata. What is one reason that a "fsck" program would still be needed, even if the WAL implementation is perfect?
- (b) Imagine a file system that uses synchronous writes (for update ordering) to protect the integrity of its metadata. What is the minimum number of disk writes that must be completed during a file create system call? Explain your answer.

- (c) Imagine that you work for a large Internet services company that has 100,000 disks in its data center. How many disk failures would you tell your boss to expect in a one year period, if each disk has a MTBF of 100 years?
- (d) Briefly describe a scenario in which Shortest-Positioning-Time-First disk scheduling could give lower service times for both requests, assuming only two disk requests are pending.
- (e) Imagine a system with an I/O workload described by a closed arrival process and zero think time. If the average service time doubles, what happens to the throughput?
- (f) You buy a disk that rotates at 6000 RPM (100 rotations per second) and has 1000 512-byte sectors on every track. When reading data from it sequentially, as fast as possible, do you expect it to provide over 50 MB/s? Explain why or why not.

Problem 2 : More short answer. [48 points]

- (a) Imagine a 5-disk disk array subsystem configured for RAID-5 with a 16 KB stripe unit size. If you were creating a log-structured file system on it, with a file system block size of 4 KB, what would be a good segment size to use? Justify your answer.
- (b) Imagine a 10-disk disk array configured to use RAID-4, with 8 data disks, one parity disk, and one spare disk. When using this disk array, Joe notices that his performance drops dramatically for about 30 minutes after a disk failure. Suggest a configuration knob that Joe might try changing in order to reduce the performance drop when he next encounters a disk failure. Explain.
- (c) Fred has constructed a perfect fsck program. When he uses it on his file system, after a crash, it reports a regular file inode with a link count of two that has no directory entries pointing to it. Describe a possible set of file system operations, performed just before the crash, that would explain how such a situation could arise.

- (d) Given an ext-2 filesystem that supports 4KB blocks, what is the largest filesize supported if inodes contain 10 direct blocks, 1 indirect block, and 1 double-indirect block? Assume block pointers are 32-bits.
- (e) Imagine a mirrored pair of disks, where each disk can service 100 I/Os per second. Given a workload that issues requests at an exponentially distributed rate w/mean 20 I/Os per second, what is the average disk response time if every request is a write?
- (f) Your fsck program finds a directory with a link count of 5. How many sub-directories should that directory have? Explain your answer.

Problem 3 : Bonus questions. [up to 2 bonus points]

(a) Which instructor is on vacation?

(b) Which school's sports teams does Greg cheer for (most strongly)?

(c) Where should Swapnil work, after he finishes his PhD?

(d) What color is Lianghong's hair?