

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. [60 points]

- (a) Large directories (i.e., those with many entries) in most file systems are very inefficient, because searching for specific entries is inefficient ($O(n)$). Give one simple approach to improving large directory efficiency without changing the on-disk directory format.

- (b) Log-structured file systems remap modified file blocks into large segments, to achieve higher write efficiency. But, doing so also requires writing a new version of each modified files' inode. Why?

- (c) Explain why a SCSI bus can only support up to 7 devices (in addition to the host's bus adapter).

- (d) What potential integrity issue arises from having a secondary index associated with a database table?

- (e) Explain why most device drivers are organized into a non-device-specific part and a collection of device-specific parts.
- (f) Most file systems allocate blocks of files near the directory that names them. Although this improves performance relative to random placement, it may not achieve the full performance of the disk when reading the set of files in a directory. Suggest two reasons why not.
- (i) First

 - (ii) Second
- (g) Many file system caches hold dirty blocks for some time, allowing updates to be coalesced or even cancelled (via deletion), but then flush them to disk. The longer they hold them, the bigger the potential performance improvement. What is the downside?
- (h) A common approach to flushing a file system cache is for a background process to periodically wake up and write out all dirty blocks. Explain why the average service time for these writes is often lower than that of cache misses.
- (i) Imagine a database table with 10000 records, each with 10 32-bit fields. Should the database system do the select first or the project first for the following query: `SELECT <attr1,attr2> FROM <table> WHERE <attr2=FOO>`? Assume 1% of records have `<attr2=FOO>`. Justify.

- (j) Disk drive firmware designers rarely expect hits in a disk drive cache for data that was recently requested by the host. Why?

- (k) Network-level flow control is used by FibreChannel switches to prevent overflowing of their internal buffers, which can happen when multiple senders try to send data to the same receiver. Give an example where flow control is needed even when there is only one sender and one receiver.

- (l) Some file systems store the data of very small files in the inode (using the space reserved for pointers to file blocks). Explain why they might be doing this.

- (m) Imagine a disk that has a single zone. Even if it exposed the number of sectors per track, Joe Smarty claims that this is not enough information to compute exactly which track held each logical block. Support Joe's assertion with a concrete example.

- (n) Why do modern disk drives normally only utilize one disk head at a time? (Rather than transferring from all heads in parallel.)

Problem 3 : Instructor trivia. [up to 2 bonus points]

(a) What is the name of Greg's younger boy?

(b) Which TA has facial hair (on purpose)?

(c) How old will Timmy be this year?

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.