

Name: _____

Instructions

There are five (5) 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!

Problem 1 : Short answer. [35 points]

- (c) Unlike client caches in most distributed file systems, no support exists in most WWW browser cache implementations for ensuring consistency.
 - (a) Why?
- (b) Explain one approach to improving the situation.
- (d) Greg has designed a super-duper deletion approach, which overwrites the data on magnetic disks such that it is truly impossible to get the data back from them. Explain why, in most environments, this would not be enough to completely eliminate all traces of the files holding his last two years of email.

Problem 2 : A little system analysis. [14 points]

Consider an e-commerce system running atop an “out-of-band” distributed storage service. The system is designed for 24/7 operation, but its front-end throttles input requests to 100 at a time, and the operators note that there are breaks between one request’s completion and the next request’s arrival. Each request requires five 1 ms interactions with the file manager and accesses to 40 file blocks. The cache hit rate is 75%, and file data is striped across five data servers such that each request has an equal probability of going to any of the disks. Each individual disk request takes 10 ms.

(a) What is the throughput of the e-commerce system?

(b) Which portion (file manager or data server) is the bottleneck? (*Show your work for full credit.*)

Problem 3 : A few disk array system design questions. [28 points]

- (a) Most high-end disk arrays include large amounts of battery-backed RAM. Explain two potential performance benefits of having such RAM.

(b) Departing briefly from disk arrays, imagine that instead of an array controller with a set of disks, clients talk directly to a set of storage servers that each have battery-backed RAM. Will the two performance benefits you listed above still work? State whether or not they will work as well and, if not, why not.

- (c) Back to disk array systems. The battery-backed RAM is also useful for integrity in disk arrays that use redundancy (e.g., mirroring or parity protection). Explain how.

 - (d) Assume, for this one part, that the RAM is **not** battery-backed. How could the RAM cache improve performance for small writes in a parity-protected disk array?

Problem 4 : File system (ext2) format and recovery. [24 points]

(a) Which of the following are *not* contained in the file system inode? Check all that apply.

- size
- access times
- inode number of parent directory
- name
- reference count

(b) Can the name of a file be determined by looking at its inode? Explain why or why not.

(c) Can the *number* of names a file has be determined by looking at its inode? Explain why or why not.

(d) Under which scenario will an inode be placed in lost+found? Check all that apply.

- One or more directories contain an entry for the inode, but the inode's reference count is zero.
- The inode's reference count is non-zero and no directories contain an entry for it.
- A file was written, but there was a power failure before the file data was written to disk.
- The inode bitmap does not show the inode as being used.

(e) Which of the following ext2 data or metadata structures are located at fixed positions on disk, that is, after formatting they never get moved. Check all that apply.

- superblock
- group descriptors
- inode bitmap
- free block bitmap
- data for a given inode (file or directory)

Problem 5 : Bonus questions. [a few bonus points]

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

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