

Stanford University
Computer Science Department
CS 240 Quiz 2
Spring 2006

May 26, 2006

This is an open-book exam. You have 50 minutes to answer 8 out of 10 questions. Write all of your answers directly on the paper. Make your answers as concise as possible. Sentence fragments ok.

NOTE: We will take off points if a correct answer also includes incorrect or irrelevant information. (I.e., don't put in everything you know in hopes of saying the correct buzzword.)

Question	Score
1 - 5	
6 - 10	
total	

Stanford University Honor Code

In accordance with both the letter and the spirit of the Honor Code, I did not cheat on this exam nor will I assist someone else cheating.

Name and Stanford ID:

Signature:

Answer eight of the following ten questions and, in a sentence or two, say *why* your answer holds (5 points each).

1. Will a guest OS be more or less susceptible to livelock when running on VMware? Let's say you are running the unmodified system in the livelock system (with screend) and that VMware can magically detect when the system is livelocked. What could it do to fix livelock?

2. When a file is created on FFS, FFS synchronously writes the directory entry pointing to the new file to disk before returning from the `creat` system call. The purported goal of doing so is to ensure that if an application later `fsync`'s the file the directory entry for the file will be on disk as well. Give an end-to-end argument against this choice.

```
3.      fd = open("foo");

        write(fd, ...)
           <----- someone else writes foo.
        read(fd, ...)
           <----- someone deletes foo
        write(fd, ...);
        close(fd);
```

What happens here on the base NFS system? For the leases system? For LBFS?

4. LBFS. Client A opens file F and starts writing. Client B opens file F and starts writing. Client B closes F. Client A closes F. When are the contents A's F? When are they B's F? When are they something else?

5. You look in an LFS segment and see a data block D immediately followed by an inode I that points to it. The LFS cleaner claims that the space used by I can be reclaimed. Does this mean D is dead as well? The cleaner claims that D is dead. Does this mean that I is as well? (Say why or why not for both.)

6. Given an intuitive estimate of the “big-O” cost of `fsck` for FFS without softupdates, for LFS? (I.e., $O(n)$, $O(n^2)$, etc where you define n .) Can you construct a specific example where FFS without softupdates obviously beats LFS by a significant amount?

7. How would network traffic change if the Map/Reduce system was modified so that map tasks created their output files in the GFS? What would the effect be on Figure 2 and Figure 3(a)? Be specific. How could this change improve performance when a map task failed?

8. Softupdates: (1) how do you use the freelist when creating a snapshot file? (2) when you write to a block on a checkpointed partition, what are the states the written block could be in and what do you do for each?

9. You take a RAID5 array with $D=4$ and mirror it (i.e., each disk in the RAID5 array has an identical copy). Give the expressions for large read, large write, large RMW, and small read, along with the intuition for why you picked these values. What is the maximum number of disks this RAID can lose before it fails?

10. You know the future. *Concisely* describe three *important* optimizations you could do (and how) for the papers covered by this quiz *besides* predicting which cache entry to evict.