

Student ID: \_\_\_\_\_

CS-736 Final: Why?  
(Fall 2007)

An Arpaci-Dusseau Exam

**Please Read All Questions Carefully!**

There are eight (8) total numbered pages

**Please put your NAME ONLY on this page, and your STUDENT ID on this and all other pages. Why? So I can grade without knowing who you are. Particularly useful for students who think I don't like them.**

**Note: I like all of you, so this is not really a problem.**

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

Student ID: \_\_\_\_\_

**This is the grading page.**

	Points	Total Possible
1		5
2		5
3		5
4		5
5		5
6		5
7		5
8		5
9		5
10		5
11		5
12		5
13		5
14		5
15		5
Total		75

Student ID: \_\_\_\_\_

Some of the big questions of the universe have no known answer:

- Why does the universe exist?
- Why are we here?
- Why do I have to take this exam?

In this exam, we focus on this simple question: why? Fortunately, in computer systems, there usually *is* an answer to this question. Thus, in this exam, you will answer why.

Why will you do this? Because you are a student; it is what you do.

**Please read each question carefully.** Then, answer why. Feel free to explain your reasoning to the level you think necessary.

Also, keep your answers **short and to the point**; there are no style points for long-winded answers (rather, quite the opposite!).





Student ID: \_\_\_\_\_

7. When there are multiple concurrent writers of a single file in AFS, it is sometimes called “last close wins”. Why?

8. GFS has “at least once” semantics for a successful record append. Why are these semantics useful for GFS?

9. In PAST, it is important that nodes that are numerically near one another in the nodeID space are not near one another in the real world. Why?

Student ID: \_\_\_\_\_

10. In “Rethink the Sync”, writes to disk are triggered in some cases by a timer, and in some cases by the generation of user-observable output (e.g., a printf). Why?

11. In the Stanford bug-finding paper, their bug-finding tool sometimes reports false positives (i.e., their tool says some code might be buggy, but actually it isn't). Why?

12. In Rx, the system sometimes changes what `malloc()` returns to the user. Why is this OK to do?

Student ID: \_\_\_\_\_

13. Processes in nucleus uses messages to communicate, not semaphores. Why?

14. FFS suggests that you write data to disk in large, 4-KB aligned chunks, instead of smaller units. Why?

15. Lottery scheduling uses randomness to implement its scheduling policy, not determinism. Why?