

Midterm I  
Fall 2003 15-712

Please write your name on the following line (and on each page of the exam):

---

This midterm is open book and open notes.

You should write brief, precise, and legible answers. Rambling brain-dump answers may be punished twice — you run out of time writing them and we run out of patience reading them. :-)

If in doubt, list any assumptions you made about the question.

There are eight questions but we will only grade **six**. List the **six** questions we should grade:

---

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

1. LFS as a database

(a) Parameterize LFS in terms of Haerder and Reuter's database taxonomy. For each parameter, circle one value and give a short (one or two sentence) justification.

Propagation Strategy: ATOMIC /  $\neg$ ATOMIC

Buffer Replacement: STEAL /  $\neg$ STEAL

End of Transaction Processing: FORCE /  $\neg$ FORCE

Checkpoint Types: TOC / TCC / ACC / FUZZY

(b) What does LFS need to undo and redo on a crash?

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

2. The eager writing disk array was designed to support transaction processing workloads. Two designers are arguing about whether adding support for track-aligned extents would improve performance for environments in which the database is frequently scanned (for report generation) in addition to its normal transaction processing activity. Give an argument for one of the two positions.

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

3. A server crash can affect a clients' ability to read and write files. For each of AFS and Pangaea, explain which files can and cannot still be read and written, and why.

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

4. (In)Secure systems

- (a) Consider a secure coprocessor such as Aegis or an IBM 4758. Suppose this hardware is running Linux, and that it has a kernel module that provides an encryption service. (Users send the encryption module a message and a key, and it returns an encrypted version.) Say this system is hacked into through a buffer overrun in the module. List three violations of Saltzer's security design principles that were exploited for this attack, and explain how.

- (b) Which class would Anderson assign to such an attack and why?

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

5. The NASD model for file services lets clients directly access objects on storage devices, but uses a “file manager” for directories and other meta-data management functions. Give and explain two reasons why it does so.

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

6. Two implementors are arguing about whether direct memory access would be a good augmentation to a River-like system. Take a position (pro or con) and explain.

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

7. A friend argued that the Pangaea architecture would be a great way to improve the Harvest proxy cache system. Give one argument for each the pro and the con.



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

8. Some systems researchers are considering extensions to Harvest to support direct memory access (à la DAFS) among nodes in the cache hierarchy. State whether you think this is good idea and explain why.