

NAME: _____ (FIRST NAME FIRST) SCORE: _____

COSC 6360 FINAL EXAMINATION DECEMBER 13, 2007

THIS EXAM IS CLOSED BOOK. YOU CAN HAVE ONE SHEET (I.E., TWO PAGES) OF NOTES. UH EXPELS CHEATERS

1. What is the purpose of the *i-node map* of a log-structured file system? (5 points) Where is it stored on disk? (5 points) How can we retrieve it when the system reboots after a crash? (5 points)

The *i-node map* of a log-structured file system contains the addresses of its *i-node* blocks. It is stored on the log along with data blocks, directory blocks and *i-node* blocks. The locations of the *i-node map* blocks are stored at a fixed location on disk each time the file system does a checkpoint.

After a crash, the file system will retrieve first the contents of its *i-node map* as it was at the time of its last checkpoint. It will then perform a roll-forward through the portion of the log after that checkpoint to retrieve more recent *i-node map* blocks.

2. What is a *stale file handle*? (5 points) Why do they occur in NFS? (5 points) What does NFS do to detect them? (5 points)

Stale file handles are file handles that point to files that have been deleted. They occur in NFS because clients cache file handles and the server does not keep track of the objects cached by the clients. NFS detects stale file handles through their generation number: a stale file handle will not have the same generation number as the file it points to.

3. How does Coda protect users against

- a) Network failures? (5 points)

By letting users access a local copy of the file "hoarded" on their workstation. ____

- b) Server failures? (5 points)

By replicating files on two—or more—servers. _____

- c) Workstation failures? (5 points)

By storing files on two—or more—central servers. _____

4. Which consistency guarantees are offered by *close-to-open consistency*? (10 points)

Whenever a client opens a file, it will always receive the most recent version of the file saved a close time by any other client.

5. Which steps do the two *journaling file systems* described by Seltzer *et al.* take to ensure that metadata blocks are always written on the log before being written at their ultimate destination on disk? (10 points)

The buffer header of each modified block in the I/O buffer identifies the last log entries describing an update to the block. Since journal blocks are written to the disk in strict FCFS order, this allows the file system to verify that all relevant log entries are written to disk before the block is flushed from the buffer.

6. What is the function of the *enode caches* in the Blue File system? (5 points) How are they organized? (5 points)

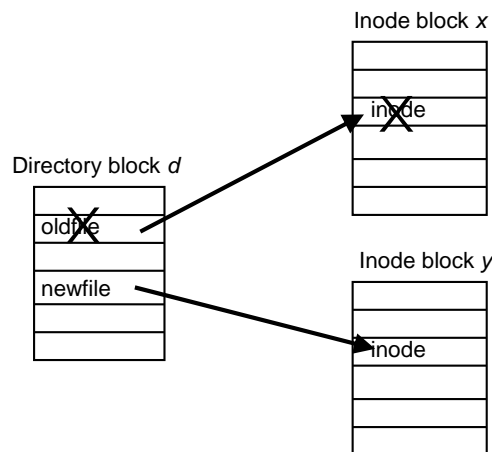
When the Blue FS has to read data it searches them on all attached devices starting with the lowest cost devices. To mitigate the impact of searching devices that do not contain the desired data, it maintains an e-node cache that records which devices hold recently used files. There node cache has affixed size. Its contents are hashed by file ID and managed by an LRU replacement policy.

7. What are the main advantage and the main disadvantage of *stateless servers*? (2×5 points)

Main advantage: stateless servers are more robust since they can be restarted at any time. (Their clients cannot tell the difference between a server that has crashed and recovered and a very slow server.)

Main disadvantage: Stateless servers cannot implement UNIX file access semantics at a reasonable cost.

8. Consider a file system using soft updates and assume that one directory block in its I/O buffer reflects the result of one file deletion and one file creation. Assuming that the i-nodes of the two files reside in two different i-node blocks, describe each step the system will take to update the disk copies of the three blocks. (3×5 points) (*Hint: treat this question as a problem.*)



- First step: write i-node block y with the new i-node to disk
- Second step: write directory block d to disk
- Third step: write i-node block x to disk

(There is no circular dependency!)