Assignment 6

(**Due date: Thursday, 12/3/2009, in class**)

Your name:	Date:

Provide **brief** answers to the following questions regarding Storage Management (Chapter 10 through Chapter 13):

The following questions are related to Chapter 10 (File System Interface):

1. Some systems automatically delete all user files when a user logs off or a job terminates, unless the user explicitly requests that they be kept; other systems keep all files unless the user explicitly deletes them. Discuss the relative merits of each approach (10.1).

2. Explain the purpose of the open() and close() operations (10.5).

3. Give an example of an application in which data in a file should be accessed in the following order (10.6):

    a. Sequentially
    b. Randomly

4. Consider a system that supports 5,000 users. Suppose that you want to allow 4,990 of these users to be able to access one file (10.8).

    a. How would you specify this protection scheme in UNIX?

5. Researchers have suggested that, instead of having an access list associated with each file (specifying which users can access the file, and how), we should have a *user control list* associated with each user (specifying which files a user can access, and how). Discuss the relative merits of these two schemes (10.9).

6. If you were creating an operating system to handle files, what would be the six basic file operations that you should implement?

7. What are common attributes that an operating system keeps track of and associates with a file?

8. What is the difference between an operating system that implements mandatory locking and one that implements advisory file locking?

The following questions are related to Chapter 11 (File System Implementation):

9. Consider a file currently consisting of 100 blocks. Assume that the file-control block (and the index block, in the case of indexed allocation) is already in memory. Calculate how many disk I/O operations are required for *contiguous, linked,* and *indexed* (single-level) allocation strategies, if, for one block, the following conditions hold. In the contiguous-allocation case, assume that here is no room to grow at the beginning but there is room to grow at the end. Also assume that the block information to be added is stored in memory (11.1).
    a. The block is added at the beginning.
    b. The block is added in the middle.
    c. The block is added at the end.
    d. The block is removed from the beginning.
    e. The block is removed from the middle.
    f. The block is removed from the end.

10. Consider a system that supports the strategies of contiguous, linked, and indexed allocation. What criteria should be used in deciding which strategy is best utilized for a particular file (11.4)?

11. Consider a file system that uses inodes to represent files. Disk blocks are 8KB in size, and a pointer to disk block requires 4 bytes. This file system has 12 direct disk blocks, as well as single, double, and triple indirect disk blocks. What is the maximum size of a file that can be stored in this file system (11.15)?

12. What do the terms "raw" and "cooked" mean when used to describe a partition?

The following questions are related to Chapter 12 (Mass Storage Structure):

13. Is disk scheduling, other than FCFS scheduling, useful in a single-user environment? Explain your answer (12.1).

14. Explain why SSTF scheduling tends to favor middle cylinders over the innermost and outermost cylinders (12.2).

15. Why is it important to balance file system I/O among the disks and controllers on a system in a multitasking environment (12.5)?

16. Suppose that a disk has 5,000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143, and the previous request was at cylinder 125. The queue of pending requests, in FIFO order, is:

    86, 1470, 913, 1774, 948, 1509, 1022, 1750, 130.

Starting from the current head position, what is the scheduling queue (list of the cylinders the drive/R-W head is to operate on: the first, second, …, last) to satisfy all the pending requests for each of the following disk-scheduling algorithms (12.16)?

a. FCFS
b. SSTF
c. SCAN
d. LOOK
e. C-SCAN
f. C-LOOK

17. Explain the terms seek time, rotational latency, and bandwidth.

18. Compare the performance of write operations achieved by a RAID Level 5 organization with that achieved by a RAID Level 1 organization (12.16).

19. Assume that you have a mixed configuration comprising disks organized as RAID Level 1 and as RAID Level 5 disks. Assume that the system has flexibility in deciding which disk organization to use for storing a particular file. Which files should be stored in the RAID Level 1 disks and which in the RAID Level 5 disks in order to optimize performance (12.27)?

==========**Important Notes**==========

- Solutions must be **typewritten**. You can use lists, bullets for the write-up (short phrases are OK; complete sentences *not* necessary).
- Put all your solutions **in the same order** as the above questions.
- Use this question paper as **cover page** and **staple them together**.
- *Electronic* submissions will be accepted ***only under*** an excusable circumstances (the above rules still apply)—in this case, put your solutions including this question paper as cover page into ***ONE SINGLE*** Word or PDF file and send it to me via email. I'll *grade your work based on this file* and send back only your grade (without corrections of your errors).
- The full score for this homework is **100** points. **You will lose 5-10 points for missing ANY ONE of the followings in your submission:**
  - No name
  - No cover page
  - Your work submitted *not* in proper order
  - Not a single Word or PDF file (if submitted via email)

==========