CS140 Operating Systems and Systems Programming
Final Exam

December 10, 2004

Name: (please print)_____

In recognition of and in the spirit of the Stanford University Honor Code, I certify that I will neither give nor receive unpermitted aid on this exam.

Signature:_____

This examination is closed notes and closed book. You may not collaborate in any manner on this exam. You have 180 minutes (3 hours) to complete the exam. Before starting, please check to make sure that you have all 19 pages.

1		**12**	
2		**13**	
3		**14**	
4		**15**	
5		**16**	
6		**17**	
7		**18**	
8			
9			
10			
11		**Total**	

Name:_____

(1) (12 points) Assume you are given a disk that is possessed. On occasion the disk has random, non-deterministic errors that might:

a) Return a different block than was requested by the software (e.g. you read block 0 but the disk returns the contents of block 1).
b) Return the correct block but one or two bits in the block flipped. (e.g. you read block 0 and it returns the contents of block 0 with the 39^{th} bit toggled).

The large majority of the time it returns the correct contents for the requested block. Describe how you as a file system designer would extend a Unix-like file system to handle such a misbehaving disk.

(2) (8 points) Given the file system from question (1), how could you handle the case in which you determined that the errors were being introduced by a malicious attacker that was trying to confuse your file system. Under this threat model the attacker can return any data for a request but must return the correct contents most of the time.

(3) (8 points) Would an operating system built with non-blocking or wait-free synchronization work correctly if run inside a virtual machine on a virtual machine monitor supporting multiple virtual machines at the same time? Justify your answer.

(4) (12 points) Describe what would happen if an operating system running in a virtual machine decided to page out a process' page and that particular page had already been paged out by the virtual machine monitor.

(5) (10 points) Assume you have two machines (A and B) sending datagram packets between one another using Internet Protocol (IP) over an internetwork. You notice from statistics being kept on the machines that machine A sent X packets to B and B received X packets from A. Machine B reports to have sent Y packets to A yet A reports to have received 4*Y packets from B. Explain how this is possible.

(6) (8 points) Your project partner builds a file system that with each file stores a list of a randomly generated 256-bit integers. In order to open a file you need to specify one of the integers in the open system call. Would this be considered an access control list or capability-based system? Justify you answer.

(7) (6 points) Explain how a buffer overrun attack works?

(8) (8 points) Recently OS researchers have proposed intelligent disks. These disks look like standard disks to the operating system yet are smart enough to parse the file system data structures on the storage to do things such as determining what blocks belong to what files and which blocks on the disk are free. Such intelligent disks could back themselves up without bothering the operating system. What attributes of modern file system make this difficult when the file system is in use by applications? Describe specific problems.

(9) (12 points) Assume you built a write-ahead logging file system on a system with two disks. One disk is used for the log and the other for the file system data.
(a) Describe what the file system must do before the space on the log can be reclaimed.
(b) Describe a situation/workload that would generated more write traffic to the log disk and than to the data disk.
(c) Describe a situation/workload that would generate more read traffic to the log disk that the data disk.

(10) (12 points) A recent piece of OS research has shown a way for an operating system to accurately predict the file access patterns of an application. It can predict if the file is
(a) read sequentially once and not access again
(b) read sequentially in a looping pattern
(b) read in neither of the above patterns
Describe how you would use this information to guide your file system buffer cache replacement policy. Describe the replacement algorithm you would suggestion for each pattern.

(11) (9 points) Explain how a RAID4 system would recover from a failure of one of the data disks.

(12) (8 points) The BSD Unix system added the fsync() system call that an application can use to make sure a file resides on disk. In other words, once the fsync() call returns changes made to the file will be safely on disk even if the system crashes. What is the minimum information that must be written to disk for fsync() to function correctly?

(13) (12 points) Fsck is the Unix file system crash recovery program. For each of the following fsck error messages, what do you believe fsck has seen in the crashed file system to generate the error message:

(A) File's inode link count is 1 should be 2.

(B) Free bitmap entry for block 3323 is 0, should be 1 (allocated).

(C) Cylinder group 4 - free block count is incorrect.

(14) (8 point) What are the advantages and disadvantages of a shortest seek time first (SSTF) disk scheduling algorithm when compared to an elevator algorithm

(15) (6 point) Would a file system that used synchronous writes of all metadata blocks see more or less benefit from a disk scheduling algorithm than one that didn't use synchronous writes of metadata?

(16) (8 point) If someone gave you the source code of a file system and asked if it used index files or extent-based allocation, what would you look for in the code to answer this question?

(17) (9 points) Describe a denial of service attack that you could perform using the virtual memory subsystem of a modern operating system.

(18) (9 points) Explain why public key cryptosystems, where keys can be freely published, are frequently deployed using trusted servers that give out keys just like are used for shared key cryptosystems where keys must be kept secret.