

Student ID: _____

**CS-736 Midterm: Fear Factor
(Fall 2002)**

Please Read All Questions Carefully!

There are eight (8) total numbered pages

Please put your NAME on this page, and your STUDENT ID on this and all other pages

Name: _____

Student ID: _____

This is the grading page.

	Points	Total Possible
Part I: Short Answers		45
Part II: Long Answers		55

Part I: Short Questions

For each of the following questions in this section, **CIRCLE ALL THAT APPLY**.

1. Which of the following could be said about the message-based communication abstraction that is used in Nucleus?
 - (a) Performance: Makes it easier to build a high-performance system
 - (b) Robustness: Makes it easier to build memory-protected higher level operating systems
 - (c) Modularity: Makes it easier to replace an existing service with a better one
 - (d) Distribution: Makes it easier to move a service to a distributed (multi-computer) setting
2. When all of the data blocks in an AutoRAID storage system are constantly being written to, and the system is exporting its maximum amount of capacity, which of the following are true?
 - (a) Most of the data is in RAID-5 storage
 - (b) Each write to the system could trigger both a block promotion and a block demotion
 - (c) The system is performing worse than a standard mirrored (RAID-1) storage system
 - (d) The system is performing worse than a standard RAID-5 storage system
3. In Pilot, a page has three flags associated with it: *referenced*, *written*, and *write-protected*. Given this, which of the following are true?
 - (a) Pilot can track (to some extent) which pages are more frequently accessed than others
 - (b) Pilot can prevent one process from reading another process's memory
 - (c) Pilot can ensure that code does not get accidentally over-written
 - (d) Pilot can track which pages are modified by a process
4. Which of the following components of a modern operating system are **not** found (in some form) in the THE system?
 - (a) Locking and synchronization
 - (b) Virtual memory paging
 - (c) Threads and concurrency
 - (d) Protected address spaces
5. To be able to virtualize the MIPS and run IRIX on top of Disco, a few changes had to be made to IRIX. Which of the following changes were made to enable correct virtualization?
 - (a) The OS code and data had to be relocated into mapped memory
 - (b) The OS idle loop had to be changed to inform Disco the OS was idling
 - (c) The OS page zeroing code had to be changed to instead request a zero'd page from Disco
 - (d) The OS bcopy had to be changed to a special Disco remapping function during network communication between virtual machines
6. Which of the following are goals in the design of exokernel?
 - (a) Avoid resource management
 - (b) Avoid multiplexing resources among competing processes
 - (c) Avoid using physical names (i.e. physical page numbers)
 - (d) Avoid visible revocation of resources

Student ID: _____

7. When the LFS cleaner is determining whether a block of a file is live, which of the following data structures might it consult?
 - (a) The inode map
 - (b) The segment summary block
 - (c) The inode of the file which the block belongs to
 - (d) The inode of the directory that contains the file
8. Which of the following were done by Chen and Bershad in their study of the impact of the structure of operating systems on memory-system performance, in order to make sure they could solely focus on structurally-induced costs and overheads?
 - (a) Slowed down the clock interrupt
 - (b) Removed the effects of policy decisions made by Ultrix and Mach
 - (c) Picked a set of workloads that stressed all components of the operating system
 - (d) Compared the run times of the workloads across the Ultrix and Mach platforms
9. The gray-box file layout detector and controller (FLDC) exploits FFS behavior to improve performance, by reading files in sorted inode-number order. Under which of the following scenarios will the FLDC likely give you little or no benefit?
 - (a) When the file system has aged significantly
 - (b) When the files that are accessed are all "small" (i.e., 4KB in size)
 - (c) When the files that are accessed are spread out across a number of directories
 - (d) When only one large file is repeatedly accessed sequentially
10. Which of the following statements can be used to describe a Hydra capability?
 - (a) Can't be forged
 - (b) Has access rights associated with it
 - (c) Can be passed from one process to another
 - (d) Is managed by the kernel
11. Cylinder groups are used in FFS to:
 - (a) Group related data blocks together
 - (b) Group related inodes together
 - (c) Group related directories together
 - (d) Increase disk efficiency
12. Which of the following is true about the original UNIX system?
 - (a) Made it easy to write, test, and share programs
 - (b) Was built to run in a memory-rich and highly CPU-powered environment
 - (c) Designers of the system used the system a lot
 - (d) Designers were shooting for the most original and innovative system they could think of

Student ID: _____

13. In comparing a standard mirrored RAID storage system (RAID level 1) of D total disks to a standard block-level striping storage system (RAID level 0) also with D disks, which of the following are **always** true?
- (a) The mirrored system appears to be half as big as the striping system to the file system
 - (b) The mirrored system tolerates more disk failures than the striping system
 - (c) The mirrored system performs worse than the striping system
 - (d) The mirrored system tolerates power outages more gracefully than the striping system
14. Vino prevents downloaded code, or *grafts*, from doing certain things within the kernel. Which of the following potential graft actions are prevented by Vino?
- (a) Changing a global scheduling policy
 - (b) Changing a thread scheduling policy for the user's threads
 - (c) Running an infinite loop
 - (d) Grabbing locks

Student ID: _____

Part II: Long Questions

1. RAIDing the File System (20 points)

In this question, you will explore issues related to running file systems efficiently on top of RAID devices.

1: Assume we are running a file system on top of RAID-5. Briefly describe RAID-5 and how it works. A picture is worth a thousand words, or at least a few points on this exam.

2: Let's say we want to run LFS on top of RAID-5. First, describe how LFS writes data to disk. What are the basic concepts and structures involved?

3: Is LFS a good file system for RAID-5, i.e., does it potentially solve some of the problems associated with RAID-5 performance? Justify.

4: To work efficiently on top of RAID-5, LFS may need to discover some parameters about the RAID system. Using *gray-box techniques*, design a simple algorithm that discovers the *stripe size* of the RAID. A loose sketch is fine (no need for a lot of details). State your assumptions. One thing you can assume is that the block size of the RAID is known.

Student ID: _____

2. The Costs of Virtualization (15 points)

This question examines Disco and the costs of virtualization. Disco virtualizes machine resources, fooling unsuspecting operating systems above into thinking they are running on the raw hardware. However, virtualization has a cost, which is exactly what we will now explore.

1: First, describe how Disco virtualizes memory. How does Disco control where memory pages get placed by the operating system?

2: Now describe how Disco virtualizes the CPU. For most instructions, what are the overheads? Which instructions are more costly to virtualize?

3: Chen and Bershad measured memory-system performance of Ultrix and Mach, to better understand the structural costs of microkernel-based operating systems. Clearly, Disco affects memory system performance as well. Discuss how Disco worsens memory-system behavior of the system. Examples would be useful.

Student ID: _____

3. Recursive Extensibility (20 points)

In this question, we study the issues surrounding systems that are designed for extensibility. Specifically, we investigate how to run one extensible OS on top of another extensible OS (recursive extensibility, you might call it).

1: Assume the base system we are running is Hydra. Describe the basic mechanisms in Hydra, concentrating on the ones that would be used to build higher-level operating systems on top of Hydra.

2: We wish to build a VINO-like system on top of Hydra. First, describe the basic mechanisms that VINO provides as an extensible OS.

3: Now, let's examine the construction of VINO on top of Hydra. How could VINO use the Hydra mechanisms to build a safe in-VINO graft execution environment? Be as specific as possible.

4: Now assume the reverse, that we wish to build Hydra on top of VINO. Can we build a Hydra environment on top of VINO?