

Name: _____

USC ID: _____

CSci 555 Final Exam

Fall 2006

Instructions:

Show all work. No electronic devices are allowed. This exam is open book, open notes. You have **115 minutes** to complete the exam.

Please prepare your answers on separate sheets of paper. You may write your answers on the sheet of paper with the question (front and back). If you need more space, please attach a separate sheet of paper to the page with the particular question. **Do NOT extend your answer on the back of the sheet for a different question, and do NOT use the same extra sheet of paper to answer more than one question.**

In particular, **each numbered questions must appear on separate pieces of paper so that the exam can be split for grading.**

Be sure to include your **name** and **USC ID** number **on each page.**

There are **100 points** in all and **3 questions.**

	Q1	Q2	Q3	Total Score
Score				

Name: _____

USC ID: _____

1. (30 points) Scalability

- a) System load (10 points) – Suggest some techniques that can be used to reduce the load on individual servers within a distributed system? Provide examples of how these techniques are used from each of the following systems: The Domain Name System, content delivery through the world wide web, remote authentication in the Kerberos system. Note that some of the systems use more than one technique.

Name: _____

USC ID: _____

- b) Identifying issues (20 points) for each of the techniques described in part (a) there are issues that must be addressed to make sure that the system functions properly (**I am interested in the properly aspect here, not the most efficiently aspect**). For each technique identify the primary issues that needs to be addressed and explain how it is addressed in each of the listed systems that uses the technique.

Name: _____

USC ID: _____

2. (35 points) Kernel Structure

For each of the operating system functions listed below list the benefits and drawbacks to placing the function in the Kernel, leaving the function to be implemented by the application, or providing the function in users space through a server (the server case includes cases where the application selects and communicates with a server, and also the case where the application calls the kernel, but the processing is redirected by the kernel to a server).

For each function, suggest the best location(s) to provide this function. If needed you can make an assumption about the scenario for which the system will be used. Justify your choice for placement of this function. There may be multiple correct answers for this last part – so long as your justification is correct.

- a. The File System (7 points)

Name: _____

USC ID: _____

b. Virtual Memory (7 points)

c. Communication (7 points)

Name: _____

USC ID: _____

d. Scheduling (7 points)

e. Security (7 points)

Name: _____

USC ID: _____

3. (35 points) Design Problem Fault Tolerance

You are designing a database system that requires significant storage and processing power. Unfortunately, you are stuck using the hardware that was ordered by the person whose job you just filled. This morning, the day after you first arrived at work, a truck arrived with 10 processors (including memory, network cards, etc), 50 disk drives, and two uninterruptible power supplies. The failure rates of the processors (including all except the disk drives and power supplies) is λ_p . The failure rates on the disk drives is λ_d , and the failure rate for the power supplies is λ_e .

a) You learned from your supervisor that the reason they let the last person go is that he designed the system so that the failure of any of the components would cause the system to stop functioning. In terms of $\lambda_p, \lambda_d, \text{and } \lambda_e$, what is the failure probability for the system as a whole. (5 points)

Name: _____

USC ID: _____

b) The highest expected load on your system could be handled by about half the processors. The largest expected dataset size that is expected is about 1/3 the capacity of the disks that arrived. Suggest a change to the structure of the system, using the components that have already arrived, that will yield better fault tolerance. In terms of $\lambda_{p,d,\text{and } e}$, what is the failure probability for the new system? (note, there are easy things and harder things you can do here, I suggest describing the easing things, generating the probability based on that approach, and then just mentioning some of the additional steps that could be taken to further improve the fault tolerance (15 points)

Name: _____

USC ID: _____

c) List some of the problems that you would need to solve or some of the assumptions you would need to make, in order to construct the system described in part b from the components that arrived this morning (things like number of network interfaces per processor, how the disks are connected to processors or the network). Discuss also any assumptions you need to make regarding detectability of failures, and describe your approach to failover (how will the failures be masked, what steps are taken when a failure occurs). (15 points)