

**Systems Qualifier Exam  
Fall, 2005**

**Name :** \_\_\_\_\_

**SID:** \_\_\_\_\_

**Score :** \_\_\_\_\_

1. List four key characteristics for the usefulness of distributed systems and explain each of them.

2. In many layered protocols, each layer has its own header. Surely it would be more efficient to have a single header at the front of each message with all the control in it than all these separate headers. Why is this not done?

3. Would it make sense to limit the number of threads in a server process?

4. Explain how DNS can be used to implement a home-based approach to locating mobile hosts.

5. List four different data-centric consistency models. What kind of consistency would you use to implement an electronic stock market?

6. Does a stateless server need to take checkpoints?

7. Rumor spreading or gossiping is a well-known technique to implement epidemic protocols. Does it guarantee eventual consistency? Explain your reason(s).

8. Constructing a concurrent server by spawning a process has some advantages and disadvantages compared to multithreaded servers. Mention a few.

9. Name two advantages and two disadvantages of using centralized servers for key management.

10. NFS does not provide a global, shared name space. Is there a way to mimic such a name space?