Architecture:
- What are the differences between user and kernel mode?
- What are the 3 names for going from user to kernel mode?
- Be able to identify from context which category a transition from user to kernel mode falls into
- Micro vs Monolithic kernel

Virtual Memory/Paging:
- What hardware is included in most systems to help support paging? How do these pieces fit together?
- Why was paging introduced over segments? What issue did it solve?
- How does paging work?
- Describe the different page replacement algorithms

Processes/Threads:
- What is the difference between a process and a thread?
- What lives in a PCB, what lives in a TCB? What does each thread have to maintain?
- Describe the differences between user threads and kernel threads
- What state can a process be in with respect to the scheduler? How can these states change?

Scheduling:
- Describe the different scheduling algorithms
- What is a multi level feedback queue?
- What do we try to maximize and minimize for effective scheduling algorithms?
- What is preemption?