

Due in class at 12:30pm on Monday 3 December 2007
typed answers preferred

1. In what way does demand paging resemble the instruction and data cache in the CPU?
2. Why is it so difficult to implement LRU page replacement, given all that's needed to do so is to record memory accesses times and maintain a sorted list?
3. What is thrashing and why is it bad?
4. What is the advantage of multilevel naming over user-based naming?
5. Why is it necessary to associate the current position in the file with a process rather than just with the file itself?