

1. **Process A** is running on a preemptive multiprocessing system. **Process B**, which is waiting for I/O, has a higher priority than **Process A**. What are the steps that must occur for **Process A** to stop running and **Process B** to start running? Be as specific as possible.

1. I/O Interrupt completing Process B's I/O (Save PC and Status)
2. Save Registers for A
3. Rearrange Ready Queue (Place B ahead of A in ready queue)
4. Restore Registers for B
5. Optional: Restore PC and Status for B

2. You have written the best sort program in the world. You have compiled it and are running it on FIRE, a LINUX system. What is the LINUX operating system doing while your sort program is running?

Nothing

3. Explain the meaning of the term "Context." What does the term "Context Switch" mean.

The context of a process is the contents of all hardware registers.

A context switch refers to the process of saving the registers of one process to memory and restoring the registers of a different process.

4. What does the term "System Call" mean? What must a program do to make a System Call happen? Give two reasons that a program might want to perform a System Call.

A system call is a request to the operating system from a process.

To make a system call happen, a process must execute a TRAP instruction.

Request I/O, Open File, Close File, Request Memory, Start Process, Kill Process, Terminate Current Process, (Many others)

5. Suppose the following jobs arrive in the following order at time zero. Give the GANTT chart and the average wait time for the following algorithms: First Come First Served, Shortest Job First. Assume that a process switch requires zero time. Calculate the average wait time for both algorithms.

Job Number	Burst Time
1	3
2	2
3	13
4	4
5	9



Average wait time = $(3+5+18+22)/5 = 48/5 = 9.6$.




Average wait time = $(2+5+9+18)/5 = 34/5 = 6.8$.

6. Suppose the following jobs enter the ready queue at the times given. Give two GANTT charts for these jobs, one for the non-preemptive Shortest Job First algorithm, and one for the preemptive Shortest Job First (aka Shortest Remaining Time First).

Job Number	Burst Time	Arrival Time
1	4	1
2	2	2
3	2	3
4	4	4
5	5	5

Non-Preemptive



Preemptive

P1	P2	P3	P1	P4	P5	
0	1	3	5	8	12	17

7. Contest winner: **A tie**, answer either *a* or *b*.
- Describe the differences between the many-to-many, many-to-one, and one-to-one relationships between kernel threads and user threads.
 - How do processes get created? Explain what creates a process and how it is done. You may use the UNIX model in your explanation.
-
- In the many-to-one relationship, threads are created in user space, but the OS has no knowledge of them. A blocking system call blocks all threads. In the one-to-one model, the OS manages all threads. In the many-to-many model, The OS manages some threads but others may be managed in user space.*
 - A process is created at the request of another process. In the UNIX or LINUX operating systems, a process is created by the “fork” system call.*