

DISCLAIMER



Everything in this document shall not, under any circumstances, hold any legal liability whatsoever. Any usage of the data and information in this document shall be solely on the responsibility of the user. This document user has to take written consent from the author.

What is preemption and example

Preemption in the Linux kernel refers to the ability to interrupt and temporarily suspend the execution of a running task (process or thread) in order to allow another task to run. Preemption ensures that tasks with higher priority or time-critical operations can be executed promptly, even if a lower-priority task is currently running.

In the Linux kernel, preemption can occur at different levels: process/thread level and interrupt level.

1. **Process/Thread Preemption:** At the process/thread level, preemption allows the kernel to preempt a running task when a higher-priority task becomes runnable. This is commonly known as "voluntary" or "cooperative" preemption. The running task voluntarily yields the CPU, allowing other tasks to run. The kernel scheduler determines the priority and order of task execution.

For example, let's say a low-priority background task is running, and a high-priority foreground task becomes ready to run. The kernel's scheduler can preempt the background task and immediately switch to executing the higher-priority foreground task, ensuring that critical tasks receive timely execution.

2. **Interrupt Preemption:** At the interrupt level, preemption allows the kernel to interrupt the execution of a running task when a higher-priority interrupt occurs. Interrupt preemption is crucial for handling time-critical events and ensuring that important interrupt handlers are executed promptly.

For example, suppose a low-priority task is executing, and a high-priority hardware interrupt, such as a timer interrupt or a network packet arrival, occurs. The kernel's interrupt handler can preempt the running task and immediately execute the interrupt handler associated with the higher-priority interrupt. Once the interrupt handler completes, the kernel resumes the execution of the original task.

Preemption in the Linux kernel plays a vital role in providing responsiveness, real-time capabilities, and efficient resource utilization. It allows the kernel to prioritize critical tasks, handle time-critical events, and prevent tasks from monopolizing system resources.

Note that the Linux kernel provides different configuration options to control the level and behavior of preemption, allowing administrators and developers to customize the system according to their specific requirements.

Preemption is particularly important in real-time and responsive systems where tasks need to be executed promptly. By preempting lower-priority tasks, the kernel can ensure that higher-priority tasks or time-critical operations are not delayed.

Let's consider an example to understand preemption better. Suppose we have two tasks in the Linux kernel: Task A and Task B. Task A has a higher priority than Task B. Without preemption, Task A would continue to execute until it voluntarily yields the CPU or completes its execution. This could potentially delay Task B from running, even if it has a higher priority.

However, with preemption enabled, the kernel can interrupt Task A and allow Task B to execute. This interruption is usually triggered by events such as a timer interrupt, I/O completion, or a higher-priority task becoming ready to run. Once the higher-priority task completes or the interrupt condition is satisfied, the kernel can switch back to Task A or allow Task A to continue its execution.

Here's a simplified example of how preemption might work in the Linux kernel:

```
void taskA(void) {
    while (1) {
        // Task A's code execution
        // ...

        // Check if Task B should be given a chance to run
        if (preemption_condition_met()) {
            // Save Task A's state and switch to Task B
            preempt_to_taskB();
        }
    }
}

void taskB(void) {
    while (1) {
        // Task B's code execution
        // ...

        // Check if Task A should be given a chance to run
        if (preemption_condition_met()) {
            // Save Task B's state and switch to Task A
            preempt_to_taskA();
        }
    }
}
```

In this example, both Task A and Task B have infinite loops representing their continuous execution. The preemption conditions are checked periodically to determine if a task switch is necessary. When a preemption condition is met, the current task's state is saved, and the kernel switches to the other task for execution.

Preemption in the Linux kernel ensures fairness and responsiveness by allowing tasks with higher priorities or time-critical operations to preempt lower-priority tasks. It plays a crucial role in maintaining system performance and meeting real-time requirements.