



## The $O(1)$ Linux Scheduler



# Introduction

---


- Scheduler: component that decides which process to run next. How to choose?
- Processes run for *time slice* units of time.
- Scheduler policy determines what runs when.
- IO-bound vs. processor bound processes.
- Priority based scheduling: runnable processes with time slice left and the highest priority always run.



# Introduction cont.

---


- Priority: based priority dynamically modified by the scheduler to fulfill scheduling policies:
  - fast process response
  - high process throughput
- Time slice: dynamic between 10 and 300ms (default 150ms). Processes can run by parts of 20ms. When the time slices are exhausted they are recalculated.



# The Scheduler

---

- Runqueue: list of runnable processors. It has 2 priority arrays: one active and one expired, swapped with pointers.
- Priority array:
  - bitmap for priorities
  - lists of processes per priority level
- Finding the highest runnable process = finding the first bit set to 1. It is *independent* on the number of processes.



# Priority Bitmap

