

Job Scheduling

You can use job scheduling to prioritize the MapReduce jobs and YARN applications that run on your MapR cluster.

The MapReduce system supports a minimum of one queue, named `default`. Hence, this parameter's value should always contain the string `default`. Some job schedulers, like the Capacity Scheduler, support multiple queues.

The default job scheduler is the Fair Scheduler, which is designed for a production environment with multiple users or groups that compete for cluster resources.

The MapR Distribution for Hadoop supports these job schedulers:

- **FIFO queue-based scheduler:** The FIFO queue scheduler runs jobs based on the order in which the jobs were submitted. You can prioritize a job by changing the value of the `mapred.job.priority` property or by calling the `setJobPriority()` method.
- **Fair Scheduler:** This is the default scheduler. The Fair Scheduler allocates a share of cluster capacity to each user over time. The design goal of the Fair Scheduler is to assign resources to jobs so that each job receives an equal share of resources over time. The Fair Scheduler enforces fair sharing within each queue. Running jobs share the queue's resources.
- **Capacity Scheduler:** The Capacity Scheduler enables users or organizations to simulate an individual hadoop cluster with FIFO scheduling for each user or organization. You can define organizations using *queues*.