

Label-based Scheduling

Label-based scheduling provides a way to allocate shared cluster resources on particular nodes in a cluster. First, you assign node labels in a text file. The node labels map to one or more nodes. Next, you can create queue labels and job labels based on the node labels.

When you run jobs, you can place them on specified nodes on a per-job basis (using a job label) or at the queue level (using a queue label). This feature is used in conjunction with schedulers, such as the [Fair Scheduler](#) or the [Capacity Scheduler](#).

To learn more about label-based scheduling for MapReduce (MRv1) jobs, see [Label-based Scheduling for MRv1 Jobs](#).

To learn more about label-based scheduling for YARN applications, such as MapReduce (MRv2) jobs, see [Label-based Scheduling for YARN Applications](#).