

Administrator Guide

Welcome to the MapR Administrator Guide! This guide is for system administrators tasked with managing MapR clusters. Topics include how to manage data by using volumes; how to monitor the cluster for performance; how to manage users and groups; how to add and remove nodes from the cluster; and more.

The focus of the Administrator Guide is managing the nodes and services that make up a cluster. The Administrator Guide does not cover the details of installing MapR software on a cluster. See [Advanced Installation Topics](#) for details on planning and installing a MapR cluster.

Click on one of the sub-sections below to get started.

- [Managing Users and Groups](#)
 - [Setting Up Email Addresses](#)
 - [Setting Up SMTP](#)
 - [Setting User Permissions](#)
 - [Setting Quotas for Users and Groups](#)
 - [Converting a Cluster from Root to Non-root User](#)
 - [Setting the Administrative User](#)
- [Managing Licenses](#)
- [Managing the Cluster](#)
 - [Managing Roles and Services](#)
 - [Adding Roles to a Node](#)
 - [Removing Roles from a Node](#)
 - [Assigning Roles to Nodes for Best Performance](#)
 - [Isolating CLDB Nodes](#)
 - [Isolating ZooKeeper Nodes](#)
 - [Task Nodes](#)
 - [TaskTracker Blacklisting](#)
 - [Setting Up MapR NFS](#)
 - [High Availability Planning](#)
 - [Starting, Stopping, and Restarting Services](#)
 - [Central Configuration](#)
 - [Changing the User for MapR Services](#)
 - [Dial Home](#)
 - [Managing Nodes](#)
 - [Adding Nodes to a Cluster](#)
 - [Performing Maintenance on a Node](#)
 - [Removing Nodes from a Cluster](#)
 - [Designating Subnets for MapR](#)
 - [Uninstalling MapR](#)
 - [Changing an IP Address](#)
 - [Managing Disks](#)
 - [Setting Up Disks for MapR](#)
 - [Handling Disk Failures](#)
 - [Working with a Logical Volume Manager](#)
 - [Shutting Down a Cluster](#)
 - [Starting Up a Cluster](#)
- [Managing Data with Volumes](#)
 - [Introduction to Volumes](#)
 - [Setting Up Topology](#)
 - [Working with Snapshots](#)
 - [Working with Mirror Volumes](#)
 - [Using Promotable Mirrors](#)
 - [Volume Upgrades in Version 5.0](#)
 - [Data Protection](#)
 - [Schedules](#)
- [Monitoring the Cluster](#)
 - [Monitoring Cluster Health](#)
 - [Monitoring Performance](#)
 - [Setting up the MapR Metrics Database](#)
 - [Service Metrics](#)
 - [Analyzing Job Metrics](#)
 - [Monitoring Node Metrics](#)
 - [Configuring Balancer Settings](#)
 - [Third-Party Monitoring Tools](#)
 - [Setting Up Alarm Notifications](#)
 - [Checking Alarms](#)
- [Cluster Resource Allocation](#)
 - [Memory Allocation for Nodes](#)
 - [Resource Allocation for Jobs and Applications](#)
- [Managing the MapReduce Mode](#)

- Configuring the Default MapReduce Mode for the Cluster
- Configuring the Default MapReduce Mode on Client Nodes
- Configuring the MapReduce Mode on the Command Line
- Configuring the MapReduce Mode with an Environment Variable
- Managing the MapReduce Mode for Ecosystem Components
- Managing Jobs and Applications
 - Job Scheduling
 - Hadoop 2.x Fair Scheduler
 - ExpressLane
 - Hadoop 2.x Capacity Scheduler
 - Label-based Scheduling
 - Submitting Jobs and Applications to the Cluster
 - Configuration Files for Jobs and Applications
 - Job and Application Logs
 - YARN Log Aggregation
 - Centralized Logging
- Maintenance Schedule
- YARN Cgroups
- JobTracker High Availability
- NodeManager Restart
- ResourceManager High Availability
 - Manual or Automatic Failover for the ResourceManager
 - Zero Configuration Failover for the ResourceManager
 - Recovery for the ResourceManager
 - ResourceManager Configuration Properties
- Performance Tuning
 - MapR-FS Performance Tuning
 - Configure Chunk Size
 - Configure NFS Write Performance
 - Increase Caching
 - MapReduce v1 Performance Tuning
 - Balance the Load Across the Cluster
 - Enable Inline Setup
 - Prevent Data Spills
 - Use Compression
 - Customize Memory Settings for MapReduce v1
 - Tune the Map Phase For MapReduce v1
 - Tune the Reduce Phase For MapReduce v1
 - Summary of Basic Guidelines For Nodes Running MapReduce v1 Jobs
 - Examples: Nodes Running MapReduce v1 Jobs
- Troubleshooting Cluster Administration