

# Managing Data with Volumes

This page describes how to organize and manage data using volumes, a unique feature of MapR clusters. This page contains the following topics:

- [Introduction to Volumes](#)
- [Creating a Volume](#)
- [Viewing a List of Volumes](#)
- [Viewing Volume Properties](#)
- [Modifying a Volume](#)
- [Mounting a Volume](#)
- [Unmounting a Volume](#)
- [Removing a Volume or Mirror](#)
- [Setting Volume Topology](#)
  - [Setting Default Volume Topology](#)
  - [Example: Setting Up CLDB-Only Nodes](#)
- [Related Topics](#)

## Introduction to Volumes

MapR provides **volumes** as a way to organize data and manage cluster performance. A volume is a logical unit that allows you to apply policies to a set of files, directories, and sub-volumes. A well-structured volume hierarchy is an essential aspect of your cluster's performance. As your cluster grows, keeping your volume hierarchy efficient maximizes your data's availability. Without a volume structure in place, your cluster's performance will be negatively affected. This section discusses fundamental volume concepts.

You can use volumes to enforce disk usage limits, set replication levels, establish ownership and accountability, and measure the cost generated by different projects or departments. Create a volume for each user, department, or project. You can mount volumes under other volumes to build a structure that reflects the needs of your organization. The volume structure defines how data is distributed across the nodes in your cluster. Create multiple small volumes with shallow paths at the top of your cluster's volume hierarchy to spread the load of access requests across the nodes.

On a cluster with an M5 license, you can create a special type of volume called a *mirror*, a local or remote read-only copy of an entire volume. Mirrors are useful for load balancing or disaster recovery.

With an M5 license, you can also create a *snapshot*, an image of a volume at a specific point in time. Snapshots are useful for rollback to a known data set. You can create snapshots and synchronize mirrors manually or using a *schedule*.

MapR lets you control and configure volumes in a number of ways:

- **Replication** - set the number of physical copies of the data, for robustness and performance
- **Topology** - restrict a volume to certain physical racks or nodes (requires M5 license and `m` permission on the volume)
- **Quota** - set a hard disk usage limit for a volume (requires M5 license)
- **Advisory Quota** - receive a notification when a volume exceeds a soft disk usage quota (requires M5 license)
- **Ownership** - set a user or group as the accounting entity for the volume
- **Permissions** - give users or groups permission to perform specified volume operations
- **File Permissions** - Unix-style read/write permissions on volumes

Volumes are stored as pieces called *containers* that contain files, directories, and other data. By default, the container size is 32 GB. Container size is controlled by the `cldb.container.size` parameter (see the `config` commands). Containers are *replicated* to protect data. There are normally three copies of each container stored on separate nodes to provide uninterrupted access to all data even if a node fails. For each volume, you can specify a *desired* replication factor and a *minimum* replication factor:

- The desired replication factor is the number of replicated copies you want to keep for normal operation and data protection. When the number of copies falls below the desired replication factor, but remains equal to or above the minimum replication factor, re-replication occurs after the timeout specified in the `cldb.fs.mark.rereplicate.sec` parameter (configurable using the `config` API). The minimum replication factor you can set is 1; the maximum is 6.
- The minimum replication factor is the smallest number of copies you want in order to adequately protect against data loss. When the replication factor falls below this minimum, re-replication occurs as aggressively as possible to restore the replication level. The minimum number allowed for the minimum replication factor is 1; the default is 2; the maximum number you can set for the minimum replication factor is 6. In all cases, the minimum replication factor cannot be greater than the replication factor.

If any containers in the CLDB volume fall below the minimum replication factor, writes are disabled until aggressive re-replication restores the

minimum level of replication. If a disk failure is detected, any data stored on the failed disk is re-replicated without regard to the timeout specified in the `cldb.fs.mark.rereplicate.sec` parameter.

## Creating a Volume

When creating a volume, the only required parameters are the volume type (normal or mirror) and the volume name. You can set the ownership, permissions, quotas, and other parameters at the time of volume creation, or use the [Volume Properties](#) dialog to set them later. If you plan to schedule snapshots or mirrors, it is useful to create a [schedule](#) ahead of time; the schedule will appear in a drop-down menu in the Volume Properties dialog.

By default, the root user and the volume creator have full control permissions on the volume. You can grant specific permissions to other users and groups:

Code	Allowed Action
dump	Dump the volume
restore	Mirror or restore the volume
m	Modify volume properties, create and delete snapshots
d	Delete a volume
fc	Full control (admin access and permission to change volume ACL)

You can create a volume using the `volume create` command, or use the following procedure to create a volume using the MapR Control System.

### To create a volume using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Click the **New Volume** button to display the [New Volume](#) dialog.
3. Use the **Volume Type** radio button at the top of the dialog to choose whether to create a standard volume, a local mirror, or a remote mirror.
4. Type a name for the volume or source volume in the **Volume Name** or **Mirror Name** field.
5. If you are creating a mirror volume:
  - Type the name of the source volume in the **Source Volume Name** field.
  - If you are creating a remote mirror volume, type the name of the cluster where the source volume resides, in the **Source Cluster Name** field.
6. You can set a mount path for the volume by typing a path in the **Mount Path** field.
7. You can specify which rack or nodes the volume will occupy by selecting a topology from the **Topology** drop-down selector.
8. You can set permissions using the fields in the **Ownership & Permissions** section:
  - a. Click **[ + Add Permission ]** to display fields for a new permission.
  - b. In the left field, type either u: and a user name, or g: and a group name.
  - c. In the right field, select permissions to grant to the user or group.
9. You can associate a standard volume with an accountable entity and set quotas in the **Usage Tracking** section:
  - a. In the **Group/User** field, select **User** or **Group** from the dropdown menu and type the user or group name in the text field.
  - b. To set an advisory quota, select the checkbox beside **Volume Advisory Quota** and type a quota (in megabytes) in the text field.
  - c. To set a quota, select the checkbox beside **Volume Quota** and type a quota (in megabytes) in the text field.
10. You can set the replication factor and choose a snapshot or mirror [schedule](#) in the Replication and Snapshot section:
  - a. Type the desired replication factor in the **Replication Factor** field. When the number of replicas drops below this threshold, the volume is re-replicated after a timeout period (configurable with the `cldb.fs.mark.rereplicate.sec` parameter using the `config API`).
  - b. Type the minimum replication factor in the **Minimum Replication** field. When the number of replicas drops below this threshold, the volume is aggressively re-replicated to bring it above the minimum replication factor.
  - c. To schedule snapshots or mirrors, select a [schedule](#) from the **Snapshot Schedule** dropdown menu or the **Mirror Update Schedule** dropdown menu respectively.
11. Click **OK** to create the volume.

## Viewing a List of Volumes

You can view all volumes using the [volume list](#) command, or view them in the MapR Control System using the following procedure.

### To view all volumes using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.

## Viewing Volume Properties

You can view volume properties using the [volume info](#) command, or use the following procedure to view them using the MapR Control System.

### To view the properties of a volume using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Display the [Volume Properties](#) dialog by clicking the volume name, or by selecting the checkbox beside the volume name, then clicking the **Properties** button.
3. After examining the volume properties, click **Close** to exit without saving changes to the volume.

## Modifying a Volume

You can modify any attributes of an existing volume, except for the following restriction:

- Mirror and normal volumes cannot be converted to the other type.

You can modify a volume using the [volume modify](#) command, or use the following procedure to modify a volume using the MapR Control System.

### To modify a volume using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Display the [Volume Properties](#) dialog by clicking the volume name, or by selecting the checkbox beside the volume name then clicking the **Properties** button.
3. Make changes to the fields. See [Creating a Volume](#) for more information about the fields.
4. After examining the volume properties, click **Modify Volume** to save changes to the volume.

## Mounting a Volume

You can mount a volume using the [volume mount](#) command, or use the following procedure to mount a volume using the MapR Control System.

### To mount a volume using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Select the checkbox beside the name of each volume you wish to mount.
3. Click the **Mount** button.

You can also mount or unmount a volume using the **Mounted** checkbox in the [Volume Properties](#) dialog. See [Modifying a Volume](#) for more information.

## Unmounting a Volume

You can unmount a volume using the [volume unmount](#) command, or use the following procedure to unmount a volume using the MapR Control System.

### To unmount a volume using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Select the checkbox beside the name of each volume you wish to unmount.

3. Click the **Unmount** button.

You can also mount or unmount a volume using the Mounted checkbox in the [Volume Properties](#) dialog. See [Modifying a Volume](#) for more information.

## Removing a Volume or Mirror

You can remove a volume using the `volume remove` command, or use the following procedure to remove a volume using the MapR Control System.

### To remove a volume or mirror using the MapR Control System:

1. In the Navigation pane, expand the **MapR-FS** group and click the [Volumes](#) view.
2. Click the checkbox next to the volume you wish to remove.
3. Click the **Remove** button to display the Remove Volume dialog.
4. In the Remove Volume dialog, click the **Remove Volume** button.

## Setting Volume Topology

You can place a volume on specific racks, nodes, or groups of nodes by setting its topology to an existing node topology.

Your node topology describes the locations of nodes and racks in a cluster. The MapR software uses node topology to determine the location of replicated copies of data. Optimally defined cluster topology results in data being replicated to separate racks, providing continued data availability in the event of rack or node failure.

Define your cluster's topology by specifying a topology for each node in the cluster. You can use topology to group nodes by rack or switch, depending on how the physical cluster is arranged and how you want MapR to place replicated data.

For more information about node topology, see [Node Topology](#).

You can set volume topology using the MapR Control System or with the `volume move` command.

### To set volume topology using the MapR Control System:

1. In the Navigation pane, expand the **MapR Data Platform** group and click the [Volumes](#) view.
2. Display the [Volume Properties](#) dialog by clicking the volume name or by selecting the checkbox beside the volume name, then clicking the **Properties** button.
3. Click **Move Volume** to display the Move Volume dialog.
4. Select a topology path that corresponds to the rack or nodes where you would like the volume to reside.
5. Click **Move Volume** to return to the Volume Properties dialog.
6. Click **Modify Volume** to save changes to the volume.

## Setting Default Volume Topology

By default, new volumes are created with a topology of `/data`. To change the default topology, use the `config save` command to change the `cldb.default.volume.topology` configuration parameter. Example:

```
maprcli config save -values "{\"cldb.default.volume.topology\":\"/data/rack02\"}"
```

After running the above command, new volumes have the volume topology `/data/rack02` by default, which could be useful to restrict new volume data to subset of the cluster.

## Example: Setting Up CLDB-Only Nodes

In a large cluster (100 nodes or more) create CLDB-only nodes to ensure high performance. This configuration also provides additional control over the placement of the CLDB data, for load balancing, fault tolerance, or high availability (HA). Setting up CLDB-only nodes involves restricting

the CLDB volume to its own topology and making sure all other volumes are on a separate topology. Because both the CLDB-only path and the non-CLDB path are children of the root topology path, new non-CLDB volumes are not guaranteed to keep off the CLDB-only nodes. To avoid this problem, set a default volume topology. See [Setting Default Volume Topology](#).

### To set up a CLDB-only node:

1. SET UP the node as usual:
  - [PREPARE](#) the node, making sure it meets the requirements.
  - [ADD](#) the MapR Repository.
2. [INSTALL](#) the following packages to the node.
  - `mapr-cldb`
  - `mapr-websserver`
  - `mapr-core`
  - `mapr-fileserver`

### To set up a volume topology that restricts the CLDB volume to specific nodes:

1. Move all CLDB nodes to a CLDB-only topology (e. g. `/cldbonly`) using the MapR Control System or the following command:  

```
maprcli node move -serverids <CLDB nodes> -topology /cldbonly
```
2. Restrict the CLDB volume to the CLDB-only topology. Use the MapR Control System or the following command:  

```
maprcli volume move -name mapr.cldb.internal -topology /cldbonly
```
3. If the CLDB volume is present on nodes not in `/cldbonly`, increase the replication factor of `mapr.cldb.internal` to create enough copies in `/cldbonly` using the MapR Control System or the following command:  

```
maprcli volume modify -name mapr.cldb.internal -replication <replication factor>
```
4. Once the volume has sufficient copies, remove the extra replicas by reducing the replication factor to the desired value using the MapR Control System or the command used in the previous step.

### To move all other volumes to a topology separate from the CLDB-only nodes:

1. Move all non-CLDB nodes to a non-CLDB topology (e. g. `/defaultRack`) using the MapR Control System or the following command:  

```
maprcli node move -serverids <all non-CLDB nodes> -topology /defaultRack
```
2. Restrict all existing volumes to the topology `/defaultRack` using the MapR Control System or the following command:  

```
maprcli volume move -name <volume> -topology /defaultRack
```

All volumes except `mapr.cluster.root` are re-replicated to the changed topology automatically.



To prevent subsequently created volumes from encroaching on the CLDB-only nodes, set a default topology that excludes the CLDB-only topology.

## Related Topics

For further information on volume-related operations, see the following topics.

- [Mirror Volumes](#)
- [Snapshots](#)
- [Schedules](#)