

Volume Upgrades in Version 5.0

Overview

MapR supports two types of volumes: read-write and mirror (read-only). The concept of promoting a mirror refers to the ability to make a read-only mirror volume into a read-write volume. The main use case for this feature is to support disaster-recovery scenarios in which a read-only mirror needs to be promoted to a read-write volume so that it can become the primary volume for data storage. You can also convert read-write volumes back to read-only mirrors.

The feature that allows mirror volumes to be promoted to read-write was introduced in Version 4.0.2; however, this feature worked only for new-format volumes (volumes that were created in new 4.0.2 clusters). Similarly, in Version 4.1, only new-format volumes that were created in 4.0.2 or 4.1 could be promoted to read-write.

However, in Version 5.0, you can upgrade old pre-4.0.2 volumes to the new format by using the `maprcli volume upgradeformat` command. After running this command on a specific volume, you can use the volume as a read-only or read-write volume by following standard procedures. New volumes that you create in Version 5.0 are created in the new format and are promotable automatically. You can continue to use old volumes as they are with no loss of functionality, but they will not be promotable.

Enabling Features for Read-Write Mirror Support

After upgrading a cluster to Version 5.0, fully enable support for read-write mirrors (promotable mirrors) by setting both of the following parameters to 1 (true). Use the `maprcli config save -values` command to set these parameters.

Parameter	Default Value	Enables/Disables
<code>dfs.feature.rwmirror.support</code>	1 (true) for Versions 5.0, 4.1, and 4.0.2 0 (false) for earlier versions	Support for the read-write mirror feature (promotable mirrors)
<code>dfs.feature.volume.upgrade</code>	0 (false)	Ability to upgrade old-type volumes so they can be used as read-write mirrors

If you prefer to allow only new volumes to be promotable, you can leave `dfs.feature.volume.upgrade` set to 0 (false). You cannot enable `dfs.feature.volume.upgrade` if `dfs.feature.rwmirror.support` is disabled.

Mirror Types

You can check the status of your volumes in terms of their "mirror type." The `maprcli volume info` command returns the following `mirrortype` values:

mirrortype	Description	Volume upgrade required (to support promotability)
0	Old-format volume, created in an earlier release and present in the cluster after an upgrade to Version 5.0	Yes, if the volume is intended for use as a read-write mirror.
1	An old-format mirror volume whose source volume is a type 0 volume (in any MapR version). These mirror volumes cannot be upgraded.	No, not allowed. The <code>maprcli volume upgradeformat</code> command returns an error for these volumes.
2	New-format mirror volume that may be promoted to read-write (no upgrade command required).	No, not needed. These volumes are already in the new format and are promotable.
3	New-format standard volume: either created new in 5.0 or upgraded in 5.0 via the <code>maprcli volume promote</code> command.	No, not needed. These volumes are already in the new format and are promotable.

To check the mirror types for your volumes in Version 5.0, run the following command:

```
maprcli volume list -columns volumename,mirrortype -json
...
{
  "volumename":"vol1999",
  "mirrortype":0
},
{
  "volumename":"volume1",
  "mirrortype":3
},
{
  "volumename":"volume2",
  "mirrortype":3
},
{
  "volumename":"volume3",
  "mirrortype":2
},
...
```

Using the maprcli volume upgradeformat Command

Use the following command to upgrade a volume from the old format to the new format:

```
maprcli volume upgradeformat [ -cluster cluster_name ] -name volumeName
```

For example, vol1999 is an old volume (mirrortype 0):

```
root@node-20:~# maprcli volume upgradeformat -name vol1999 -json
{
  "timestamp":1433465410299,
  "timeofday":"2015-06-04 05:50:10.299 GMT-0700",
  "status":"OK",
  "total":0,
  "data":[
  ]
}
```

After the upgradeformat command has been run, the volume's mirror type is now 3:

```
root@node-20:~# maprcli volume info -name vol1999 -json | grep mirrortype
  "mirrortype":3,
```

The upgradeformat operation takes some time to complete, depending on the number of containers, but the volume remains accessible for reads and writes during the operation. The mirror type does not get updated until the operation is complete.

Do not attempt to mirror data to a volume that is in the process of being upgraded. A mirroring operation will hang and must be aborted if it attempts to mirror from an old source volume that is in the process of being upgraded to the new type.

The system does not provide any progress information during the upgradeformat operation.

The upgradeformat command returns an error in the following situations:

- If the volume is already a promotable (new-format) volume:

```
mapr@node-20:~$ maprcli volume create -name volumel
...
mapr@node-20:~$ maprcli volume upgradeformat -name volumel
ERROR (22) - Volume upgradeformat failed: volumel is already in new type.
Invalid argument
```

- If the volume is an old-format mirror volume. You cannot upgrade a mirror volume that was created in the old format.



Upgraded source volumes must mirror to new-type mirror volumes. If you upgrade to Version 5.0 and upgrade your source volumes, you cannot continue to use old-type mirror volumes for mirroring. You must create new mirror volumes in Version 5.0. However, a maprcli command (maprcli volume mirror start) that attempts to mirror to an old mirror volume from a new source volume will not fail at the CLI prompt. The maprcli interface does not validate volume types. You need to monitor the attempted mirror operation itself to detect any errors.