

volume create

Creates a volume. Permissions required: cv or fc on the volume

Syntax

CLI	<pre>maprcli volume create -name <volume name> [-advisoryquota <advisory quota>] [-ae <accounting entity>] [-aetype <accounting entity type>] [-allowgrant true false] [-allowinherit true false] [-auditenabled true false] [-cluster <cluster>] [-coalesce interval in mins] [-createparent 0 1] [-group <list of group:allowMask>] [-inherit <volume name>] [-localvolumehost <localvolumehost>] [-localvolumeport <localvolumeport>] [-maxinodesalarmthreshold <maxinodesalarmthreshold>] [-minreplication <minimum replication factor>] [-mirrorschedule <mirror schedule ID>] (4.0.2 only) [-mount 0 1] [-nsminreplication <minimum replication factor>] [-nsreplication <replication factor>] [-path <mount path>] [-quota <quota>] [-readonly <read-only status>] [-replication <replication factor>] [-replicationtype <type>] [-rereplicationtimeoutsec <seconds>] [-rootdirperms <root directory permissions>] [-schedule <ID>] [-source <source>] [-topology <topology>] [-type rw mirror] (4.0.2) or [-type 0 1] (4.0.1) [-user <list of user:allowMask>]</pre>
REST	<pre>http[s]://<host>:<port>/rest/volume/create?<parameters></pre>

Parameters

Parameter	Description
advisoryquota	The advisory quota for the volume as integer plus unit. Example: quota=500G; Units: B, K, M, G, T, P
ae	The accounting entity that owns the volume.
aetype	The type of accounting entity: <ul style="list-style-type: none">• 0=user• 1=group
allowgrant	Specifies whether a parent volume grants permission for a child volume to inherit its properties. True or false. Default: false.

allowinherit	Specifies whether a new volume inherit sproperties from the parent mount point volume. True or false. Default: true.
auditenabled	<p>Specifies whether to turn auditing on for the volume. If auditing is also enabled at the cluster level with the <code>maprccli audit data</code> command, setting this value to <code>true</code> causes auditing to start for any directories, files, or tables that are already enabled for auditing. If none are yet so enabled, subsequently enabling auditing on any of them causes auditing of them to start.</p> <p>You must have the <code>fc</code> permission on the cluster to use this parameter. See <code>acl</code> for details about this permission.</p>
cluster	The cluster on which to create the volume.
coalesce	<p>The interval of time during which READ, WRITE, or GETATTR operations on one file from one client IP address are logged only once, if auditing is enabled.</p> <p>For example, suppose that a client application reads a single file three times in 6 minutes, so that there is one read at 0 minutes, another at 3 minutes, and a final read at 6 minutes. If the coalesce interval is at least 6 minutes, then only the first read operation is logged. However, if the interval is between 4 minutes, then only the first and third read operations are logged. If the interval is 2 minutes, all three read operations are logged.</p> <p>The default value is 60 minutes. Setting this field to a larger number helps prevent audit logs from growing quickly.</p>
createparent	<p>Specifies whether or not to create a parent directory to hold the volume link:</p> <ul style="list-style-type: none"> • 0 = Do not create a parent directory. • 1 = Create a parent directory.
group	Space-separated list of <code>group:permission</code> pairs.
inherit	Specifies the name of the volume from which the new volume will inherit properties. When <code>inherit</code> is explicitly specified, <code>allowgrant</code> is not required.
localvolumehost	The local volume host.
localvolumeport	The local volume port. Default: 5660
maxinodesalarmthreshold	Threshold for the <code>INODES_EXCEEDED</code> alarm.
minreplication	The minimum replication level. Default: 2 When the replication factor falls below this minimum, re-replication occurs as aggressively as possible to restore the replication level. 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.
mirrorschedule	The ID of a mirror schedule. If a mirror schedule ID is provided, then the mirror volume will automatically sync with its source volume on the specified schedule.
mount	Specifies whether the volume is mounted at creation time.
name	The name of the volume to create.
nsminreplication	The minimum replication level for the namespace container. Default: 2 When the replication factor falls below this minimum, re-replication occurs as aggressively as possible to restore the replication level. 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. When enabled, the CLDB manages the namespace container replication separate from the data container replication. This capability is used when you have low volume replication but want to have higher namespace replication. Note: This parameter must be the same or larger than the equivalent data replication parameter, <code>minreplication</code> .
nsreplication	The desired namespace container replication level. Default: 3. 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 <code>cldb.fs.mark.rereplicate.sec</code> parameter. This timeout is the time given for a down node to come back online. After this timeout period, the CLDB takes action to restore the replication factor. When enabled, the CLDB manages the namespace container replication separate from the data container replication. This capability is used when you have low volume replication but want to have higher namespace replication. Note: This parameter must be the same or larger than the equivalent data replication parameter, <code>replication</code> .
path	The path at which to mount the volume.
quota	<p>The quota for the volume as <code>integer plus unit</code>. Example: <code>quota=500G</code>; Units: B, K, M, G, T, P</p> <p>Do not use two-letter abbreviations for quota units, such as GB and MB.</p>

readonly	<p>Specifies whether or not the volume is read-only:</p> <ul style="list-style-type: none"> • 0 = Volume is read/write. • 1 = Volume is read-only.
replication	<p>The desired replication level. Default: 3 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 <code>cl db.fs.mark.rereplicate.sec</code> parameter. Note that this timeout is the time given for a down node to come back online. After this timeout period, the CLDB takes action to restore the replication factor.</p>
replicationtype	<p>The desired replication type. You can specify <code>low_latency</code> (star replication) or <code>high_throughput</code> (chain replication). The default setting is <code>high_throughput</code>.</p>
rereplicationtimeoutsec	<p>Timeout (in seconds) before attempting re-replication of replica containers. This volume property defines the timeout period until CLDB starts re-replicating the containers on the node of the volume when CLDB stops receiving a heartbeat from the node.</p> <p>When a node is down, CLDB gives the node an hour to come back online before it takes any action for the containers on this node. This parameter can be set on volumes to reduce the default 1 hour to a shorter time period. This option is provided mainly for local volumes, <code>repl=1</code> so that when an mfs is down, CLDB can give up quickly and decide that the container has no master. This forces the TT to give up on local containers and take appropriate recovery action of deleting the mapped volume and creating another one.</p>
rootdirperms	<p>Permissions on the volume root directory.</p>
schedule	<p>The ID of a schedule. Use the schedule list command to find the ID of the named schedule you want to apply to the volume.</p> <p>In 4.0.1, if a schedule ID is provided, it applies to a mirror volume. The mirror volume will automatically sync with its source volume on the specified schedule. Use the schedule list command to find the ID of the named schedule you wish to apply to the volume.</p> <p>In 4.0.2 and later releases, if a schedule ID is provided, it applies to either a mirror volume or a standard volume. A standard volume will automatically create snapshots on the specified schedule. A mirror volume will sync with its source volume on the specified schedule. If a mirror volume is made into a standard volume, the <code>-schedule</code> option is used as the snapshot schedule for the standard volume.</p>
source	<p>For mirror volumes, the source volume to mirror, in the format <code><source volume>@<cluster></code> (Required when creating a mirror volume).</p>
topology	<p>The rack path to the volume.</p>
type	<p>The type of volume to create.</p> <p>For 4.0.2 and later releases, these values are accepted:</p> <ul style="list-style-type: none"> • <code>mirror</code> - standard mirror (read-only) volume (promotable to standard read-write volume) • <code>rw</code> - standard (read-write) volume (convertible to standard mirror volume) • <code>0</code> - standard (read-write) volume (for backward compatibility) • <code>1</code> - non-convertible mirror (read-only) volume (for backward compatibility) <p>For 4.0.1 and earlier releases, these values are accepted:</p> <ul style="list-style-type: none"> • <code>0</code> - standard (read-write) volume (this volume is non-convertible) • <code>1</code> - standard mirror (read-only) volume (this volume is non-convertible)
user	<p>Space-separated list of <code>user:permission</code> pairs.</p>

Examples

Create the volume "test-volume" mounted at "/test/test-volume"

CLI	<pre>maprcli volume create -name test-volume -path /test/test-volume -type rw</pre>
REST	<pre>https://rln1.sj.us:8443/rest/volume/create?name=test-volume&path=/test/test-volume&type=rw</pre>

Create Volume with a Quota and an Advisory Quota

This example creates a volume with the following parameters:

- advisoryquota: 100M
- name: volumename
- path: /volumepath
- quota: 500M
- replication: 3
- schedule: 2
- topology: /East Coast
- type: rw

CLI	<pre>maprcli volume create -name volumename -path /volumepath -advisoryquota 100M -quota 500M -replicatio</pre>
REST	<pre>https://r1n1.sj.us:8443/rest/volume/create?advisoryquota=100M&name=volumename&path=/volumepath&quota</pre>

Create the mirror volume "test-volume.mirror" from source volume "test-volume" and mount at "/test/test-volume-mirror"

CLI	<pre>maprcli volume create -name test-volume.mirror -source test-volume@src-cluster-name -path /test/test</pre>
REST	<pre>https://r1n1.sj.us:8443/rest/volume/create?name=test-volume.mirror&source=test-volume@src-cluster-nam</pre>

Create volumes that inherit from a parent volume

When creating and mounting a volume, the location of the mount path is specified by the path parameter. Volumes can be mounted via the web console, the `maprcli` commands, or the REST commands. The `maprcli` commands include `volume create -path` command and the `maprcli volume mount -path` command if the volume was previously created. Sub-volumes (children) can inherit properties from their parent volume.

In the following example, a parent volume and two (2) child volumes are create where the child volume inherit properties from the parent. When the `inherit` flag is explicitly used, the `allowgrant` parameter for the parent volume is not required.

- For child volumes, c1 and c2, inheritance is explicit because the `inherit` option is specified. Thus, p1.c1 and p1.c2 volumes will inherit all properties from volume p1 (note that p1 is not a parent of p1.c1) regardless of whether the `allowgrant` option is set on p1 or not. In this case, there is an explicit inheritance ant the `allowgrant` flag is ignored and volume properties are inherited.
- For the child volume, c3, inheritance is implicit. Meaning, the child volume, p1.c3, inherits all properties from the parent volume, p1, only if the `allowgrant` option is set on p1.

CLI	<pre>maprcli volume create -name p1 -path /p1 maprcli volume create -name p1.c1 -inherit p1 maprcli volume create -name p1.c2 -path /p1/c2 -inherit p1 maprcli volume create -name p1.c3 -path /p1/c3</pre>
REST	<pre>https://10.10.82.26:8443/rest/volume/create?name=p1&path=%2Fp1 https://10.10.82.26:8443/rest/volume/create?name=p1.c1&inherit=p1 https://10.10.82.26:8443/rest/volume/create?name=p1.c2&path=%2Fp1%2Fc2?inhe https://10.10.82.26:8443/rest/volume/create?name=p1.c3&path=%2Fp1%2Fc3</pre>

In the following example, the p1.child volume normally inherits from the p1 parent volume properties because p1.child is mounted under p1 and the allowgrant option is set to true on the parent volume. However, if the child volume doesn't want to inherit properties, then set the allowinherit option to false (default: true).

CLI	<pre>maprcli volume create -name p1 -path /p1 -allowgrant true maprcli volume create -name p1.child -path /p1/p1.child -allowinherit false</pre>
REST	<pre>https://10.10.82.26:8443/rest/volume/create?name=p1&path=%2Fp1&allowgrant=true https://10.10.82.26:8443/rest/volume/create?name=p1.child&path=%2Fp1%2Fp1.child&allowinherit=false</pre>

Create a volume with namespace container replicas

CLI	<pre>maprcli volume create -name testVol -nsminreplication 1 -nsreplication 2</pre>
REST	<pre>https://10.10.82.26:8443/rest/volume/create?name=testVol&nsminreplication=1&nsreplication=2</pre>