

For example:

```
# gluster volume rebalance test-volume stop
Node          Rebalanced    size scanned failures skipped status   run time
              -files
-----
localhost      18 595.5KB    204      0      49 stopped   0:0:5
10.70.37.202    0 0Bytes      0      0      0 stopped   0:0:5
volume rebalance: test-volume: success: rebalance process may be in the
middle of a file migration.
The process will be fully stopped once the migration of the file is
complete.
Please check rebalance process for completion before doing any further brick
related tasks on the volume.
```

## 11.8. Setting up Shared Storage Volume

Features like Snapshot Scheduler, NFS Ganesha and geo-replication require a shared storage to be available across all nodes of the cluster. A gluster volume named **gluster\_shared\_storage** is made available for this purpose, and is facilitated by the following volume set option.

```
cluster.enable-shared-storage
```

This option accepts the following two values:

» **enable**

When the volume set option is enabled, a gluster volume named **gluster\_shared\_storage** is created in the cluster, and is mounted at **/var/run/gluster/shared\_storage** on all the nodes in the cluster.



### Note

- This option cannot be enabled if there is only one node present in the cluster, or if only one node is online in the cluster.
- The volume created is either a replica 2, or a replica 3 volume. This depends on the number of nodes which are online in the cluster at the time of enabling this option and each of these nodes will have one brick participating in the volume. The brick path participating in the volume is **/var/lib/glusterd/ss\_brick**.
- The mount entry is also added to **/etc/fstab** as part of **enable**.
- Before enabling this feature make sure that there is no volume named **gluster\_shared\_storage** in the cluster. This volume name is reserved for internal use only

After successfully setting up the shared storage volume, when a new node is added to the cluster, the shared storage is not mounted automatically on this node. Neither is the **/etc/fstab** entry added for the shared storage on this node. To make use of shared storage on this node, execute the following commands:

```
# mount -t glusterfs <local node's ip>:gluster_shared_storage
/var/run/gluster/shared_storage
# cp /etc/fstab /var/run/gluster/fstab.tmp
# echo "<local node's ip>:/gluster_shared_storage
/var/run/gluster/shared_storage/ glusterfs defaults 0 0" >> /etc/fstab
```

» **disable**

When the volume set option is disabled, the **gluster\_shared\_storage** volume is unmounted on all the nodes in the cluster, and then the volume is deleted. The mount entry from **/etc/fstab** as part of **disable** is also removed.

For example:

```
# gluster volume set all cluster.enable-shared-storage enable
volume set: success
```



### Important

After creating a cluster execute the following command on all nodes present in the cluster:

```
systemctl enable glusterfssharedstorage.service
```

This is applicable for Red Hat Enterprise Linux 7 only.

## 11.9. Stopping Volumes

To stop a volume, use the following command:

```
# gluster volume stop VOLNAME
```

For example, to stop test-volume:

```
# gluster volume stop test-volume
Stopping volume will make its data inaccessible. Do you want to continue?
(y/n) y
Stopping volume test-volume has been successful
```

## 11.10. Deleting Volumes



### Important

Volumes must be unmounted and stopped before you can delete them. Ensure that you also remove entries relating to this volume from the **/etc/fstab** file after the volume has been deleted.

To delete a volume, use the following command:

```
# gluster volume delete VOLNAME
```

For example, to delete test-volume: