

Red Hat Storage 2.1

2.1 Release Notes

Release Notes for Red Hat Storage



Pavithra Srinivasan

Shalaka Harne

Divya Muntimadugu

Red Hat Storage 2.1 2.1 Release Notes

Release Notes for Red Hat Storage

Edition 1

| | | |
|--------|---------------------|--|
| Author | Pavithra Srinivasan | psriniva@redhat.com |
| Author | Shalaka Harne | sharne@redhat.com |
| Author | Divya Muntimadugu | divya@redhat.com |

Copyright © 2013 Red Hat Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

Java® is a registered trademark of Oracle and/or its affiliates.

XFS® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

All other trademarks are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Storage 2.1.

| | |
|---|-----------|
| Preface | v |
| 1. Introduction | 1 |
| 2. What is New in this Release? | 3 |
| 3. Known Issues | 5 |
| 4. Technology Previews | 15 |
| 4.1. Red Hat Storage Console | 15 |
| 4.2. Striped Volumes | 15 |
| 4.3. Distributed-Striped Volumes | 15 |
| 4.4. Distributed-Striped-Replicated Volumes | 15 |
| 4.5. Striped-Replicated Volumes | 16 |
| 4.6. Replicated Volumes with Replica Count greater than 2 | 16 |
| 4.7. Support for RDMA over Infiniband | 16 |
| 4.8. Stopping Remove Brick Operation | 16 |
| 4.9. Read-only Volume | 16 |
| A. Revision History | 17 |

Preface

The Red Hat Storage Release Notes document lists the changes (that is, new features and known issues) in this release. It also contains a complete list of all currently available Technology Preview features.

Should you require information regarding the Red Hat Storage life cycle, refer to <https://access.redhat.com/support/policy/updates/rhs/>.

Introduction

Red Hat Storage is a software only, scale-out storage solution that provides flexible and agile unstructured data storage for the enterprise. Red Hat Storage provides new opportunities to unify data storage and infrastructure, increase performance, and improve availability and manageability in order to meet a broader set of an organization's storage challenges and needs.

GlusterFS, a key building block of Red Hat Storage, is based on a stackable user space design and can deliver exceptional performance for diverse workloads. GlusterFS aggregates various storage servers over network interconnects into one large parallel network file system. The POSIX compatible GlusterFS servers, which use XFS file system format to store data on disks, can be accessed using industry standard access protocols including Network File System (NFS) and Server Message Block SMB (also known as CIFS).

Red Hat Storage can be deployed in the private cloud or data center using Red Hat Storage Server for On-premise. Red Hat Storage can be installed on commodity servers and storage hardware resulting in a powerful, massively scalable, and highly available NAS environment. Additionally, Red Hat Storage can be deployed in the public cloud using Red Hat Storage Server for Public Cloud, for example, within the Amazon Web Services (AWS) cloud. It delivers all the features and functionality possible in a private cloud or datacenter to the public cloud by providing massively scalable and high available NAS in the cloud.

Red Hat Storage Server for On-Premise

Red Hat Storage Server for On-Premise enables enterprises to treat physical storage as a virtualized, scalable, and centrally managed pool of storage by using commodity server and storage hardware.

Red Hat Storage Server for Public Cloud

Red Hat Storage Server for Public Cloud packages GlusterFS as an Amazon Machine Image (AMI) for deploying scalable NAS in the AWS public cloud. This powerful storage server provides a highly available, scalable, virtualized, and centrally managed pool of storage for Amazon users.

What is New in this Release?

This chapter describes the key features added to Red Hat Storage 2.1.

- **Distributed Geo-replication**

With this release, the Geo-replication process of glusterFS is distributed, and synchronizes the local changes of each brick (on each node) parallelly to the remote slave node. The consistency guarantee of glusterFS for synchronizing the data is more reliable in this release.

For more information, refer to *Chapter 11. Managing Geo-replication* in the *Red Hat Storage 2.1 Administration Guide*.

- **SMB Enhancements**

The performance of the read and write operations in Red Hat Storage has improved.

For more information, refer to *Section 9.3. SMB* of *Chapter 9. Accessing Data - Setting Up Clients* in the *Red Hat Storage 2.1 Administration Guide*.

- **Red Hat Storage as the storage platform for Red Hat OpenStack**

Red Hat Storage is a software based distributed technology that is scalable and highly available as software only. Red Hat Storage can be deployed in the cloud or datacenter using Red Hat Storage Server.

Red Hat OpenStack provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

Red Hat Storage integration with Red Hat OpenStack is hardened and validated by Red Hat and is best suited to serve as the storage platform for Red Hat OpenStack.

For more information, refer to the *Red Hat Storage 2.1 Installation and Configuration Guide*.

- **Object Store**

Object Store technology enables enterprises to adopt and deploy cloud storage solutions. It allows users to access and modify data as objects from a REST interface along with the ability to access and modify files from NAS interfaces.

Object Storage has now been rebased on OpenStack Grizzly version.

For more information, refer to *Chapter 18. Managing Object Store* in the *Red Hat Storage 2.1 Administration Guide*.

- **Red Hat Storage Console (Technical Preview)**

- **Import Cluster feature in Red Hat Storage Console**

With this release, you can import a Red Hat Storage cluster and all the hosts belonging to the cluster into the Red Hat Storage Console.

For more information, refer to *Chapter 3. Managing Cluster* in the *Red Hat Storage 2.1 Console Administration Guide*.

- **Gluster Sync-Hosts,Volume,Brick**

Chapter 2. What is New in this Release?

The Gluster Sync periodically fetches the latest cluster configuration from GlusterFS and synchronizes the same with the engine database. This can be performed through the Red Hat Storage Console.

For more information, refer to *Chapter 4. Managing Storage Servers* in the *Red Hat Storage 2.1 Console Administration Guide*.

- **Gluster Hooks Management**

Gluster Hooks are volume life cycle extensions. They can be managed from the Red Hat Storage Console. The content of the hook can be viewed if the content type of the hook is **Text**.

For more information, refer to *Chapter 6. Managing Gluster Hooks* in the *Red Hat Storage 2.1 Console Administration Guide*.

- **Detailed Information of a Brick**

The advanced details of a particular brick of the volume can be viewed through the Red Hat Storage Console. The **Advanced View** displays the details of the brick which is divided into four parts namely General, Clients, Memory Statistics, and Memory Pools.

For more information, refer to *Chapter 5. Managing Volumes* in the *Red Hat Storage 2.1 Console Administration Guide*.

- **Optimize Volume for Virt Store**

Red Hat Storage Volumes can be optimized for virtualization through the Red Hat Storage Console.

For more information, refer to *Chapter 5. Managing Volumes* in the *Red Hat Storage 2.1 Console Administration Guide*.

Known Issues

This chapter provides a list of known issues at the time of release.

- **Issues related to Red Hat Enterprise Virtualization and Red Hat Storage Integration**

- If the Red Hat Storage server nodes and the Red Hat Enterprise Virtualization Hypervisors are present in the same data center, the servers of both types are listed for selection when you create a virtual machine or add a storage domain. Red Hat recommends that you create a separate data center for the Red Hat Storage server nodes.

- BZ# [867236](https://bugzilla.redhat.com/show_bug.cgi?id=867236)¹

While deleting a virtual machine using the Red Hat Enterprise Virtualization Manager, the virtual machine is deleted but remains in the actual storage. This consumes unnecessary storage.

Workaround: Delete the virtual machine manually using the command line interface. The virtual image file is deleted to free the space.

- BZ# [918032](https://bugzilla.redhat.com/show_bug.cgi?id=918032)²

In this release, the **direct-io-mode=enable** mount option does not work on the Hypervisor.

- BZ# [920791](https://bugzilla.redhat.com/show_bug.cgi?id=920791)³ and BZ# [920530](https://bugzilla.redhat.com/show_bug.cgi?id=920530)⁴

In a plain distributed hash table (DHT), there is no assurance of data availability leading to the unavailability of virtual machines. This may result in disruption of the cluster.

For a high availability requirement, it is recommended that you use distributed-replicate volumes on the Hypervisors.

- BZ# [979901](https://bugzilla.redhat.com/show_bug.cgi?id=979901)⁵

Virtual machines may experience very slow performance when a rebalance operation is initiated on the storage volume. This scenario is observed when the load on storage servers are extremely high. Hence, it is recommended to run the rebalance operation when the load is low.

- BZ# [856121](https://bugzilla.redhat.com/show_bug.cgi?id=856121)⁶

When a volume starts, a **.glusterfs** directory is created in the back-end export directory. When a **remove-brick** command is performed, it only changes the volume configuration to remove the brick and stale data is present in back-end export directory.

Workaround: Run this command on the Red Hat Storage Server node to delete the stale data.

```
rm -rf /export-dir
```

- BZ# [866908](https://bugzilla.redhat.com/show_bug.cgi?id=866908)⁷

¹ https://bugzilla.redhat.com/show_bug.cgi?id=867236

² https://bugzilla.redhat.com/show_bug.cgi?id=918032

³ https://bugzilla.redhat.com/show_bug.cgi?id=920791

⁴ https://bugzilla.redhat.com/show_bug.cgi?id=920530

⁵ https://bugzilla.redhat.com/show_bug.cgi?id=979901

⁶ https://bugzilla.redhat.com/show_bug.cgi?id=856121

⁷ https://bugzilla.redhat.com/show_bug.cgi?id=866908

The `gluster volume heal <volname> info` command gives stale entries in its output in a few scenarios.

Workaround: Execute the command after 10 minutes. This removes the entries from internal data structures and the command does not display stale entries.

- **Issues related to Red Hat Storage Console**

- BZ# [922572](https://bugzilla.redhat.com/show_bug.cgi?id=922572)⁸

Jboss application is updated after the Red Hat Console is installed causing an HTTP 500 error while accessing the console through the web interface.

Workaround: Edit the `standalone.xml` located in `jbossas/standalone/configuration/` by removing the `<user-name>` tag from security element under data source element.

- BZ# [905440](https://bugzilla.redhat.com/show_bug.cgi?id=905440)⁹

Due to a bug in JBoss modules (<https://issues.jboss.org/browse/MODULES-105>), the Red Hat Storage Console may not work after the latest patches are applied.

Workaround: After every yum update run this command:

```
# find /usr/share/jbossas/modules -name '*.jar.index' -delete
```

And then restart the jbossas service.

- BZ# [916981](https://bugzilla.redhat.com/show_bug.cgi?id=916981)¹⁰

In this release, VDSM supports the functionality of cluster compatibility level 3.1. Hence, only a Red Hat Storage 2.0 server with a compatibility level 3.1 data center can be added to a cluster using Red Hat Enterprise Virtualization Manager.

- BZ# [916095](https://bugzilla.redhat.com/show_bug.cgi?id=916095)¹¹

When a server is added to a cluster through the Red Hat Storage Console using the IP address and consequently if the server is added to the cluster again using the hostname; the action does not fail right away. Instead, the Console attempts to perform the installation and then fails. The **newly-added** host goes to the **Install Failed** state.

- BZ# [989477](https://bugzilla.redhat.com/show_bug.cgi?id=989477)¹²

The `restore.sh` script fails to restore the engine database when run with a user other than `postgres`. You can run the `restore.sh` script only with `-u postgres` option.

- BZ# [972581](https://bugzilla.redhat.com/show_bug.cgi?id=972581)¹³

⁸ https://bugzilla.redhat.com/show_bug.cgi?id=922572

⁹ https://bugzilla.redhat.com/show_bug.cgi?id=905440

¹⁰ https://bugzilla.redhat.com/show_bug.cgi?id=916981

¹¹ https://bugzilla.redhat.com/show_bug.cgi?id=916095

¹² https://bugzilla.redhat.com/show_bug.cgi?id=989477

¹³ https://bugzilla.redhat.com/show_bug.cgi?id=972581

The **list events --show-all** command and the **show event <id>** command raises a Python error with the datetime object. This renders the list events and show event CLI commands unusable.

- BZ# [990108](#)¹⁴

Resetting the **user .cifs** option using the **Create Volume** operation on the **Volume Options** tab on the Red Hat Storage Console reports **Error while executing action Reset Gluster Volume Options: Volume reset failed.**

- BZ# [970581](#)¹⁵

When attempting to select a volume option from the **Volume Option** drop down list, the list collapses before you make a selection.

Workaround: Click **Volume Option** again to make a selection.

- BZ# [989382](#)¹⁶

No errors are reported when you start the `ovirt-engine-notifier`. There is no notification that the `ovirt-engine-notifier` started successfully.

Workaround: Check the status of the service using the command:

```
#service ovirt-engine-notifier status
```

- BZ# [1007751](#)¹⁷

During the installation of the `rhsc-setup` RPM, the following benign warnings are seen because the `ovirt` user-id and user-group is not created.

- warning: user ovirt does not exist - using root
- warning: group ovirt does not exist - using root

- **Issues related to the Red Hat Storage Console Command Line Interface:**

- BZ# [928926](#)¹⁸

When you create a cluster, both the `glusterFS` service and `virt` service get enabled on the server. An HTTP error message must be displayed and there should be a restriction on creating a cluster with both the services enabled at the same time.

- **Issues related to Rebalancing Volumes:**

- Rebalance does not happen if the bricks are down.

While running `rebalance`, ensure that all the bricks are in the operating or connected state.

- BZ# [960910](#)¹⁹

¹⁴ https://bugzilla.redhat.com/show_bug.cgi?id=990108

¹⁵ https://bugzilla.redhat.com/show_bug.cgi?id=970581

¹⁶ https://bugzilla.redhat.com/show_bug.cgi?id=989382

¹⁷ https://bugzilla.redhat.com/show_bug.cgi?id=1007751

¹⁸ https://bugzilla.redhat.com/show_bug.cgi?id=928926

¹⁹ https://bugzilla.redhat.com/show_bug.cgi?id=960910

After rebalancing a volume, if you run `rm -rf` command on the mount point to remove all of the content from the current working directory recursively without prompting, you may get *Directory not Empty* error message.

- BZ# [862618](https://bugzilla.redhat.com/show_bug.cgi?id=862618)²⁰

After completion of the rebalance operation, there may be mismatch of failure counts between the **gluster volume rebalance status** output and the rebalance log files.

- BZ# [980081](https://bugzilla.redhat.com/show_bug.cgi?id=980081)²¹

Applications that run on the Red Hat Storage server might fail and encounter an I/O error when the **add-brick** command is executed continuously while the I/O operation is in progress on the mount point.

Workaround : Run rebalance each time you run an **add-brick** command.

- BZ# [987327](https://bugzilla.redhat.com/show_bug.cgi?id=987327)²²

If the user performs a rename operation on some files while the rebalance operation is in progress, some of those files might not be visible on the mount point after the rebalance operation is complete.

- **Issues related to Self-heal**

- BZ# [877895](https://bugzilla.redhat.com/show_bug.cgi?id=877895)²³

When one of the bricks in a replicate volume is offline, the `ls -lR` command from the mount point reports *Transport end point not connected*.

When one of the two bricks under replication goes down, the entries are created on the other brick. The Automatic File Replication translator remembers that the directory that is down contains stale data. If the brick that is up is killed before the self-heal happens on that directory, operations like `readdir()` fail.

- BZ# [972021](https://bugzilla.redhat.com/show_bug.cgi?id=972021)²⁴

In certain cases due to a race condition of network connectivity, opening a file before the completion of the self-heal process, leads to the file having stale data.

- BZ# [852294](https://bugzilla.redhat.com/show_bug.cgi?id=852294)²⁵

If the number of files which need to be self-healed is large, the GlusterFS CLI reports *Operation failed* for the command **gluster volume heal vol info**.

- BZ# [920970](https://bugzilla.redhat.com/show_bug.cgi?id=920970)²⁶

If the **gluster volume heal info** command hangs, subsequent commands fail for the next 10 minutes due to the cluster-wide lock time out.

²⁰ https://bugzilla.redhat.com/show_bug.cgi?id=862618

²¹ https://bugzilla.redhat.com/show_bug.cgi?id=980081

²² https://bugzilla.redhat.com/show_bug.cgi?id=987327

²³ https://bugzilla.redhat.com/show_bug.cgi?id=877895

²⁴ https://bugzilla.redhat.com/show_bug.cgi?id=972021

²⁵ https://bugzilla.redhat.com/show_bug.cgi?id=852294

²⁶ https://bugzilla.redhat.com/show_bug.cgi?id=920970

- **Issues related to replace-brick operation**

- Even though the **replace-brick** status command displays *Migration complete*, all the data would not have been migrated to the destination brick. It is *strongly* recommended that you get cautious when performing the **replace-brick** operation.
- The replace-brick operation will not be successful if either the source or the destination brick goes down.
- After the **gluster volume replace-brick VOLNAME Brick New-Brick commit** command is executed, the file system operations on that particular volume, which are in transit, fail.
- After a replace-brick operation, the stat information is different on the NFS mount and the FUSE mount. This happens due to internal time stamp changes when the **replace-brick** operation is performed.

- **Issues related to Directory Quota:**

The Directory Quota feature is under Technology Preview.

- BZ# [1001453](https://bugzilla.redhat.com/show_bug.cgi?id=1001453)²⁷

Quota limit enforcement may be incorrect if a file is expanded to a larger size using the truncate() system call.

- BZ# [1002885](https://bugzilla.redhat.com/show_bug.cgi?id=1002885)²⁸

Adding a brick to a volume does not enforce quota limit on the files that are on the newly added bricks.

- BZ# [1003755](https://bugzilla.redhat.com/show_bug.cgi?id=1003755)²⁹

Directory Quota feature does not work well with hardlinks. With a directory that has Quota limit set, the disk usage seen with the **du -hs <directory>** command and the disk usage seen with the **gluster volume quota VOLNAME list <directory>** command may differ. It is recommended that applications writing to a volume with directory quotas enabled, do not use hardlinks.

- **Issues related to NFS**

- After you restart the NFS server, the unlock within the grace-period feature may fail and the locks help previously may not be reclaimed.
- **fcntl** locking (NLM) does not work over IPv6.
- You cannot perform an NFS mount on a machine on which glusterFS-NFS process is already running. This is because glusterFS-NFS has already registered the NLM port with portmapper. When you start the NFS client, it also tries to register NLM with the portmapper and fails.

Workaround: Use the NFS mount **-o nolock** option to mount.

²⁷ https://bugzilla.redhat.com/show_bug.cgi?id=1001453

²⁸ https://bugzilla.redhat.com/show_bug.cgi?id=1002885

²⁹ https://bugzilla.redhat.com/show_bug.cgi?id=1003755

- If the NFS client is behind a NAT (Network Address Translation) router or a firewall, the locking behavior is unpredictable. The current implementation of NLM assumes that Network Address Translation of the client's IP does not happen.
- **nfs.mount-udp** option is disabled by default. You must enable it if you want to use posix-locks on Solaris when using NFS to mount a glusterFS volume.
- If you enable the **nfs.mount-udp** option, while mounting a subdirectory (exported using the **nfs.export-dir** option) on Linux, you must mount using the **-o proto=tcp** option. UDP is not supported for subdirectory mounts on the GlusterNFS server.
- For NLM to function properly, you must ensure that all of the servers and clients have resolvable hostnames. That is, servers must be able to resolve client names and clients must be able to resolve server hostnames.

- BZ# [973078](#)³⁰

For a distributed or a distributed-replicated volume, in the case of an NFS mount, if the brick or sub-volume is down, then any attempt to create, access, or modify the file which is either hashed or hashed and cached on the sub-volume that is down gives an *I/O error* instead of a *Transport endpoint is not connected* error.

- **Issues related to Object Store**

- The GET and PUT commands fail on large files while using Unified File and Object Storage.

Workaround: You must set the `node_timeout=60` variable in the proxy, container, and the object server configuration files.

- BZ# [985862](#)³¹

When you try to copy a file larger than that of the brick size, an HTTP error 503 is returned.

Workaround: Increase the amount of storage available in the corresponding volume and retry.

- **Issues related to distributed Geo-replication**

- BZ# [980910](#)³²

Changes in the meta data file on the master are not propagated to the slave.

- BZ# [984813](#)³³

The files which were removed on the master volume when Geo-replication was stopped, will not be removed from the slave, when Geo-replication restarts.

- BZ# [984591](#)³⁴

After stopping a Geo-replication session, if the files synced to the slave volume are renamed then when Geo-replication starts again, the renamed files are treated anew, (without considering the

³⁰ https://bugzilla.redhat.com/show_bug.cgi?id=973078

³¹ https://bugzilla.redhat.com/show_bug.cgi?id=985862

³² https://bugzilla.redhat.com/show_bug.cgi?id=980910

³³ https://bugzilla.redhat.com/show_bug.cgi?id=984813

³⁴ https://bugzilla.redhat.com/show_bug.cgi?id=984591

renaming) and synced on to the slave volumes again. For example, if 100 files were renamed, you would find 200 files on the slave side.

- BZ# [984603](#)³⁵

After stopping a Geo-replication session, hardlinks are created to point to the files. These hardlinks are treated as new files when the Geo-replication session is restarted. Therefore, the total disk consumption on the slave volume is greater than master volume disk consumption.

- BZ# [987929](#)³⁶

While the rebalance process is in progress, starting or stopping a Geo-replication session results in some files not get synced to the slave volumes. a Geo-replication sync process is in progress, running the rebalance command causes the Geo-replication sync process to stop. As a result, some files do not get synced to the slave volumes.

- BZ# [1000948](#)³⁷

If there are tens of millions of files on the master volume and you start a Geo-replication session, it takes a very long time for you to observe the updates on the slave mount point.

- **Issues related to glusterFS**

- BZ# [877988](#)³⁸

Creating a hard link fails with the **File Exist** message even though the hard link file does not exist. This occurs when a replicate-brick in a replication pair is not connected properly.

Workaround: Execute this command to avoid this issue:

```
gluster volume set VOLNAME stat-prefetch off
```

- BZ# [979861](#)³⁹

Although the glusterd service is alive, the **gluster** command reports glusterd as non-operational.

Workaround: There are two ways to solve this:

Edit `/etc/glusterfs/glusterd.vol` to contain this line:

```
option rpc-auth-allow-insecure on
```

Or

Reduce the `tcp_fin_timeout` from default 60 seconds to 1 second

The `tcp_fin_timeout` variable tells the kernel how long to keep sockets in the state FIN-WAIT-2 if you were the one closing the socket.

³⁵ https://bugzilla.redhat.com/show_bug.cgi?id=984603

³⁶ https://bugzilla.redhat.com/show_bug.cgi?id=987929

³⁷ https://bugzilla.redhat.com/show_bug.cgi?id=1000948

³⁸ https://bugzilla.redhat.com/show_bug.cgi?id=877988

³⁹ https://bugzilla.redhat.com/show_bug.cgi?id=979861

- BZ# [986090](https://bugzilla.redhat.com/show_bug.cgi?id=986090)⁴⁰

Currently, the Red Hat Storage server has issues with mixed usage of hostnames, IPs and FQDNs to refer to a peer. If a peer has been probed using its hostname but IPs are used during add-brick, the operation may fail. It is recommended to use the same address for all the operations, that is, during peer probe, volume creation, and adding/removing bricks. It is preferable if the address is correctly resolvable to a FQDN.

- BZ# [882769](https://bugzilla.redhat.com/show_bug.cgi?id=882769)⁴¹

When a glusterFS volume is started, by default the NFS and SMB server processes are also start automatically. The simultaneous use of SMB or NFS protocols to access the same volume is not supported.

Workaround: You must ensure that the same volume is accessed either using SMB or NFS protocols.

- BZ# [852293](https://bugzilla.redhat.com/show_bug.cgi?id=852293)⁴²

The management daemon does not have a rollback mechanism to revert any action that may have succeeded on some nodes and failed on the those that do not have the brick's parent directory. For example, setting the volume-id extended attribute may fail on some nodes and succeed on others. Because of this, the subsequent attempts to recreate the volume using the same bricks may fail with the error *<brickname> or a prefix of it is already part of a volume.*

Workaround:

1. You can either remove the brick directories or remove the glusterfs-related extended attributes.
2. Try creating the volume again.

- BZ# [977492](https://bugzilla.redhat.com/show_bug.cgi?id=977492)⁴³

If the NFS client machine has more than 8 GB RAM and if the virtual memory subsystem is set with the default value of `vm.dirty_ratio` and `vm.dirty_background_ratio`, the NFS client caches a huge amount of write-data before committing it to the GlusterFS NFS server. The GlusterFS NFS server does not handle huge I/O bursts, it slows down and eventually stops.

Workaround: Set the virtual memory parameters to increase the NFS COMMIT frequency to avoid huge I/O bursts. The suggested values are:

```
vm.dirty_background_bytes=32768000
```

```
vm.dirty_bytes=65536000
```

Issues related to POSIX ACLs:

- Mounting glusterFS with `-o acl` can negatively impact directory read performance. Commands like recursive directory listing can be slower than normal.

⁴⁰ https://bugzilla.redhat.com/show_bug.cgi?id=986090

⁴¹ https://bugzilla.redhat.com/show_bug.cgi?id=882769

⁴² https://bugzilla.redhat.com/show_bug.cgi?id=852293

⁴³ https://bugzilla.redhat.com/show_bug.cgi?id=977492

-
- When POSIX ACLs are set and multiple NFS clients are used, there could be inconsistency in the way ACLs are applied due to attribute caching in NFS. For a consistent view of POSIX ACLs in a multiple client setup, use the **-o noac** option on the NFS mount to disable attribute caching. Note that disabling the attribute caching option could lead to a performance impact on the operations involving the attributes.

- **Issues related to GlusterFS Samba**

- BZ# [994990](https://bugzilla.redhat.com/show_bug.cgi?id=994990)⁴⁴

When the same file is accessed concurrently by multiple users for reading and writing. The users trying to write to the same file will not be able to complete the write operation because of the lock not being available.

Workaround: To avoid the issue, execute the command:

```
gluster volume set <VOLNAME> storage.batch-fsync-delay-usec 0
```

- **General issues**

- If files and directories have different GFIDs on different back-ends, the glusterFS client may hang or display errors.

Contact Red Hat Support for more information on this issue.

- BZ# [865672](https://bugzilla.redhat.com/show_bug.cgi?id=865672)⁴⁵

Changing a volume from one-brick to multiple bricks (add-brick operation) is not supported. The volume operations on the volume may fail due to impact of add brick operation on the volume configuration.

It is recommended that the volume is started with at least two bricks to avoid this issue.

- BZ# [839213](https://bugzilla.redhat.com/show_bug.cgi?id=839213)⁴⁶

A volume deleted in the absence of one of the peers is not removed from the cluster's list of volumes. This is due to the import logic of peers that rejoin the cluster. The import logic is not capable of differentiating between deleted and added volumes in the absence of the other (conflicting) peers.

Work Around : Manually detect it by analyzing the CLI **cmd** logs to get the cluster view of the volumes that must have been present. If any volume is not listed, use the **volume-sync** command to reconcile the volumes in the cluster.

- BZ# [920002](https://bugzilla.redhat.com/show_bug.cgi?id=920002)⁴⁷

The POSIX compliance tests fail in certain cases on RHEL 5.9 due to mismatched timestamps on FUSE mounts. These tests pass on all the other RHEL 5.x and RHEL 6.x clients.

⁴⁴ https://bugzilla.redhat.com/show_bug.cgi?id=994990

⁴⁵ https://bugzilla.redhat.com/show_bug.cgi?id=865672

⁴⁶ https://bugzilla.redhat.com/show_bug.cgi?id=839213

⁴⁷ https://bugzilla.redhat.com/show_bug.cgi?id=920002

- BZ# [916834](https://bugzilla.redhat.com/show_bug.cgi?id=916834)⁴⁸

The **quick-read** translator returns stale file handles for certain patterns of file access. When running the *dbench* application on the mount point, a *dbench: read fails on handle 10030* message is displayed.

Work Around: Use the command below to avoid the issue:

```
gluster volume set VOLNAME quick-read off
```

⁴⁸ https://bugzilla.redhat.com/show_bug.cgi?id=916834

Technology Previews

This chapter provides a list of all available Technology Preview features in Red Hat Storage 2.1.

Technology Preview features are currently not supported under Red Hat Storage subscription services, may not be functionally complete, and are generally not suitable for production use. However, these features are included as a customer convenience and to provide the feature with wider exposure.

Customers may find these features useful in a non-production environment. Customers are also free to provide feedback and functionality suggestions for a Technology Preview feature before it becomes fully supported. Errata will be provided for high-severity security issues.

During the development of a Technology Preview feature, additional components may become available to the public for testing. It is the intention of Red Hat to fully support Technology Preview features in a future release.

4.1. Red Hat Storage Console

Red Hat Storage Console is a powerful and simple web based Graphical User Interface for managing a Red Hat Storage 2.1 environment. It helps Storage Administrators to easily create and manage multiple storage pools. This includes features like elastically expanding or shrinking a cluster, creating and managing volumes.

For more information, refer to *Red Hat Storage 2.1 Console Administration Guide*.

4.2. Striped Volumes

Striped volumes stripes data across bricks in the volume. Use striped volumes only in high concurrency environments accessing very large files is critical.

For more information, refer to section *Creating Striped Volumes* in the *Red Hat Storage 2.1 Administration Guide*.

4.3. Distributed-Striped Volumes

The distributed striped volumes stripe data across two or more nodes in the trusted storage pool. Use distributed striped volumes to scale storage and to access very large files during critical operations in high concurrency environments.

For more information, refer to section *Creating Distributed Striped Volumes* in the *Red Hat Storage 2.1 Administration Guide*.

4.4. Distributed-Striped-Replicated Volumes

Distributed striped replicated volumes distributes striped data across replicated bricks in a trusted storage pool. Use distributed striped replicated volumes in highly concurrent environments where there is parallel access of very large files and performance is critical. Configuration of this volume type is supported only for Map Reduce workloads.

For more information, refer to the section *Creating Distributed Striped Replicated Volumes* in the *Red Hat Storage 2.1 Administration Guide*.

4.5. Striped-Replicated Volumes

The striped replicated volumes stripe data across replicated bricks in a trusted storage pool. Use striped replicated volumes in highly concurrent environments where there is parallel access of very large files and performance is critical. In this release, configuration of this volume type is supported only for Map Reduce workloads.

For more information, refer to the section *Creating Striped Replicated Volumes* in the *Red Hat Storage 2.1 Administration Guide*.

4.6. Replicated Volumes with Replica Count greater than 2

The replicated volumes create copies of files across multiple bricks in the volume. You can use replicated volumes in environments where high-availability and high-reliability are critical. Creating replicated volumes with replica count more than 2 is under technology preview.

For more information, refer to the section *Creating Replicated Volumes* in the *Red Hat Storage 2.1 Administration Guide*.

4.7. Support for RDMA over Infiniband

Red Hat Storage support for RDMA over Infiniband is a technology preview feature.

4.8. Stopping Remove Brick Operation

You can cancel a remove-brick operation. After executing a remove-brick operation, you can choose to stop the remove-brick operation by executing the **stop** command. The files that are already migrated during remove-brick operation, is not migrated back to the same brick.

For more information, refer to the section *Stopping Remove Brick Operation* in the *Red Hat Storage 2.1 Administration Guide*.

4.9. Read-only Volume

Red Hat Storage enables you to mount volumes with read-only permission. While mounting the client, you can mount a volume as read-only and also make the entire volume as read-only, which applies for all the clients using the **volume set** command.

Appendix A. Revision History

Revision 2.1-10 **Thurs September 4 2013** **Pavithra Srinivasan** psriniva@redhat.com

Updated the What's New and Known Issues sections for the GA release.

Revision 2.1-4 **Thurs August 1 2013** **Pavithra Srinivasan** psriniva@redhat.com

Updated the What's New and Technology Preview sections for RC release.

Revision 2.1-3 **Wed July 31 2013** **Pavithra Srinivasan** psriniva@redhat.com

Updated the What's New and Technology Preview sections for RC release.

Revision 2.1-2 **Wed July 31 2013** **Pavithra Srinivasan** psriniva@redhat.com

Updated the What's New and Technology Preview sections for RC release.

Revision 2.1-1 **Tue July 30 2013** **Pavithra Srinivasan** psriniva@redhat.com

Updated the known issues section with bugs for RC release. Also updated the What's New section.

Revision 2.1-0 **Thu July 25 2013** **Shalaka Harne** sharne@redhat.com

Created the What's New section and updated the Technical Preview section.

Revision 1.1-0 **Mon Jun 3 2013** **Divya Muntimadugu** divya@redhat.com

Draft version of the document for the 2.1 release.

