

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's New?	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

Red Hat Storage Release Notes documents the changes (that is, bugs fixed, enhancements added, and known issues found) in this release. This document 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 Common Internet File System (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's New?

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.

- **SMB Enhancements**

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

- **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 hosts belonging to the cluster into the Red Hat Storage Console.

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

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.

- **Gluster Hooks Management**

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

- **Viewing advanced details for 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.

- **Optimize Volume for Virt Store**

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

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**

- A split-brain is observed on the virtual machine image files, if the source brick for self-heal operation goes off-line before the self-heal is completed.
- The scatter-gather I/O patch significantly improves the performance of the virtual machines hosted on the fuse-enabled Red Hat Storage Servers. This patch has been integrated into the Red Hat Enterprise Linux 6.3.z kernel-2.6.32-279.22.1.el6. Ensure that this kernel patch is updated on your system.
- sosreport generation throws errors.

Workaround: Bootstrap the Red Hat Storage server nodes using Red Hat Enterprise Virtualization Manager or manually create `/etc/vdsm/vdsm.conf` on all Red Hat Storage server nodes. You must ensure to delete the `/etc/vdsm/vdsm.conf`, if the nodes are bootstrapped later.

- If 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 during the creation of a virtual machine or adding a storage domain. Red Hat recommends that you create a separate data center for Red Hat Storage server nodes.
- BZ# [867236](https://bugzilla.redhat.com/show_bug.cgi?id=867236)¹

While deleting a virtual machine using Red Hat Enterprise Virtualization Manager results in the virtual machine being deleted but not from the actual storage. This consumes unnecessary storage.

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

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

In this release, `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 which leads to unavailability of virtual machines. This may result in disruption of the cluster.

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

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

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

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 rebalance operation when the load is low.

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

When a volume starts, `.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 `rm -rf /export-dir` command on the Red Hat Storage Server node to delete the stale data.

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

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

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

- BZ# [871727](https://bugzilla.redhat.com/show_bug.cgi?id=871727)⁸

If a VM image is hosted on a pure replicate volume, the VM image goes to a non-responsive state when one of the nodes having a replica brick is powered off.

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

- BZ# [922572](https://bugzilla.redhat.com/show_bug.cgi?id=922572)⁹

Jboss application updated after Red Hat Console is installed causes 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)¹⁰

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

Workaround: After every yum update, run `# find /usr/share/jbossas/modules -name '*.jar.index'` command. Restart the jbossas service to successfully log into the Red Hat Storage Console.

- BZ# [916981](https://bugzilla.redhat.com/show_bug.cgi?id=916981)¹¹

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

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

⁸ https://bugzilla.redhat.com/show_bug.cgi?id=871727

⁹ https://bugzilla.redhat.com/show_bug.cgi?id=922572

¹⁰ https://bugzilla.redhat.com/show_bug.cgi?id=905440

¹¹ https://bugzilla.redhat.com/show_bug.cgi?id=916981

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

- BZ#[916095](https://bugzilla.redhat.com/show_bug.cgi?id=916095)¹²

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 it fails. The **newly-added** host goes to the **Install Failed** state.

- BZ# [989477](https://bugzilla.redhat.com/show_bug.cgi?id=989477)¹³

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 option `-u postgres`.

- BZ# [972581](https://bugzilla.redhat.com/show_bug.cgi?id=972581)¹⁴

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](https://bugzilla.redhat.com/show_bug.cgi?id=990108)¹⁵

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

- BZ# [970581](https://bugzilla.redhat.com/show_bug.cgi?id=970581)¹⁶

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.

- **Issues related to Rebalancing Volumes:**

- Rebalance does not happen if the bricks are down.

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

- After rebalancing a volume, if you run `rm -rf` command at the mount point to remove all contents of the current working directory recursively without prompting, you may get *Directory not Empty* error message.
- Rebalance operation fails to migrate data in the distributed-striped volume.
- Rebalance operation on geo-replicated volumes in the slave can lead to data inconsistency on the slave till files are updated on master.
- If one of the bricks in a replicated volume goes off-line, the rebalance operation can fail.

¹² https://bugzilla.redhat.com/show_bug.cgi?id=916095

¹³ https://bugzilla.redhat.com/show_bug.cgi?id=989477

¹⁴ https://bugzilla.redhat.com/show_bug.cgi?id=972581

¹⁵ https://bugzilla.redhat.com/show_bug.cgi?id=990108

¹⁶ https://bugzilla.redhat.com/show_bug.cgi?id=970581

Workaround: Restart the rebalance operation after the brick is on-line.

- 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.
- While the rebalance operation is in progress, the **file already exists** errors in the brick log file can be ignored.
- BZ# [980081](#)¹⁷

Applications that run on Red Hat Storage 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](#)¹⁸
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**

- On the NFS mount, few files can go missing till the **self-heal** daemon heals all the files for a short duration during migration.

Workaround: Run a full file system crawl after the migration and the files can be viewed on the NFS mount.

- If the source brick of a self-heal operation in a replicated volume goes off-line even after successful completion of self-healing on virtual machine files, the virtual machine goes into a pause state.
- If entries are created in a directory while one of the bricks in a replicated volume is down and the source of the brick for self-heal goes off-line before the self-heal operation is complete on that directory, then operations like **readdir** on that directory fails until self-heal operation is completed.
- BZ# [877895](#)¹⁹

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](#)²⁰

¹⁷ https://bugzilla.redhat.com/show_bug.cgi?id=980081

¹⁸ https://bugzilla.redhat.com/show_bug.cgi?id=987327

¹⁹ https://bugzilla.redhat.com/show_bug.cgi?id=877895

²⁰ https://bugzilla.redhat.com/show_bug.cgi?id=972021

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)²¹

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

- BZ# [920970](https://bugzilla.redhat.com/show_bug.cgi?id=920970)²²

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

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

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

- **Issues related to Directory Quota:**

Directory Quota is a technology preview feature.

- Some writes appear to pass even though the quota limit is exceeded (write returns success). This is because they could be cached in write-behind. However disk-space does not exceed the quota limit, since quota does not allow writes to the backend. Hence, it is recommended that applications check the return value of close call.
- Excessive logging while deleting files when *Quota* or *gsync-indexing* options are enabled.
- When a user changes directory on which the administrator is setting the limit, even though the command succeeds, the new limit value is applicable to all the users except for those users' who have changed into that particular directory. The old limit value will be applicable until the user has moved out of that directory.
- BZ# [848253](https://bugzilla.redhat.com/show_bug.cgi?id=848253)²³

When quota limit is set on a distributed volume and a brick goes down while I/O is in progress, the effective *quota limit* can exceed as the distribute translator is not aware of the contribution from the offline brick.

²¹ https://bugzilla.redhat.com/show_bug.cgi?id=852294

²² https://bugzilla.redhat.com/show_bug.cgi?id=920970

²³ https://bugzilla.redhat.com/show_bug.cgi?id=848253

It is recommended that you use a replicated volume if quota usage is critical when a node goes down. In a replicate volume, quota limit is maintained even during single brick failure.

- Rename operation (that is, removing old path and creating new path) requires additional disk space equal to the file size. This is because, during rename, it subtracts the size on old path after rename operation is performed, but it checks whether quota limit is exceeded on parents of new file before rename operation.
- With striped volumes, the directory quota feature is not available.
- When renaming a file, if the available free size is less than the size of the file, quota displays "Disk limit exceeded" error without renaming the file.
- When you set quota limit on a directory that contains ',' (comma) in its name, then quota implementation may not work.
- **Issues related to POSIX ACLs:**
 - Even though POSIX ACLs are set on the file or directory, the + (plus) sign in the file permissions is not displayed. This results in performance optimization and will be fixed in a future release.
 - When glusterFS is mounted with **-o acl**, directory read performance can be bad. Commands like recursive directory listing can be slower than normal.
 - 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 **-o noac** option on NFS mount to switch off attribute caching. This could have a performance impact on operations involving attributes.
- **Issues related to NFS**
 - After you restart the NFS server, the unlock within the grace-period may fail and previously held locks may not be reclaimed.
 - **fcntl** locking (NLM) does not work over IPv6.
 - You cannot perform NFS mount on a machine on which glusterFS-NFS process is already running. This is because glusterFS-NFS has already registered NLM port with portmapper. So, when you start NFS client, it also tries to register NLM with portmapper and fails.

Workaround: You must use NFS mount **-o nolock** option to mount.
 - If the NFS client is behind a firewall such as NAT (Network Address Translation) router, the locking behavior is unpredictable. The current implementation of NLM assumes there are no NAT happening to client's IP.
 - **nfs.mount-udp** option is disabled by default. You must enable it if you want to use posix-locks on solaris when you NFS mount the gluster volume.
 - If you enable **nfs.mount-udp** option, while mounting a subdirectory (exported using **nfs.export-dir** option) on Linux, you must mount using **-o proto=tcp** option. UDP is not supported for subdirectory mounts on GlusterNFS server.
 - For NLM to function properly, you must ensure that all 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](https://bugzilla.redhat.com/show_bug.cgi?id=973078)²⁴

For a distributed or a distributed-replicated volume, in 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 Unified File and Object Storage**

- GET and PUT commands fails on the large files while using Unified File and Object Storage.

Workaround: You must ensure to add `node_timeout=60` variable in proxy, container, and object server configuration files.

- In this release, Object Expiration feature of Swift-1.4.8 version is not supported.

- **Issues related to geo-replication**

- Geo-replication uses `rsync` to sync files from master to slave, but `rsync` does not sync mknod and pipe files.

- If you start a geo-replication session on a volume that has a replace-brick operation in progress, there is a possibility of data loss on the slave.

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

- BZ# [980910](https://bugzilla.redhat.com/show_bug.cgi?id=980910)²⁵

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

- BZ# [989192](https://bugzilla.redhat.com/show_bug.cgi?id=989192)²⁶

With the current build the Geo-replication failover-failback process is not supported.

- **Minor known missing feature:**

- locks - mandatory locking is not supported.

- **Issues related to glusterFS**

- BZ# 989435

On the gluster volume, setting the Access Control List does not work as expected due to the caching of metadata in client side. As a result of this, the Access Control List setting with `setfacl` command is not visible from `getfacl` command.

Workaround : Disable `stat-prefetch` by running the command:

```
gluster volume set <VOL> stat-prefetch disable.
```

Result : After disabling `stat-prefetch`, the Access Control Lists work fine on the gluster mount point.

²⁴ https://bugzilla.redhat.com/show_bug.cgi?id=973078

²⁵ https://bugzilla.redhat.com/show_bug.cgi?id=980910

²⁶ https://bugzilla.redhat.com/show_bug.cgi?id=989192

- BZ# [877988](https://bugzilla.redhat.com/show_bug.cgi?id=877988)²⁷

Creating hard link fails with the **File Exist** message even though the hard link file does not exist. This issue is due to replicate bricks are not connected properly.

Workaround: Execute **gluster volume set VOLNAME stat-prefetch off** command to avoid this issue.

- BZ# [979861](https://bugzilla.redhat.com/show_bug.cgi?id=979861)²⁸

Although the glusterd service is alive, **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 fin_timeout.

tcp_fin_timeout from default 60seconds to 1 second

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

- BZ# [986090](https://bugzilla.redhat.com/show_bug.cgi?id=986090)²⁹

Currently, Red Hat Storage 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# 928784

Automatic File Replication algorithm assumes a certain order for committing the change log extended attributes and corresponding data operations. When there is a system crash or hard reset, the back-end file system (XFS) re-orders the operations and possibly lose the data operations. This results in the system going to an inconsistent or irrecoverable state and in extreme cases, it might lead to data corruption.

- BZ# [882769](https://bugzilla.redhat.com/show_bug.cgi?id=882769)³⁰

When a Gluster volume is started, by default both NFS and CIFS server processes are also start automatically. The simultaneous access of CIFS and NFS to the same volume is not supported.

Workaround: Ensure that the same volume is not accessed using CIFS and NFS protocol simultaneously.

²⁷ https://bugzilla.redhat.com/show_bug.cgi?id=877988

²⁸ https://bugzilla.redhat.com/show_bug.cgi?id=979861

²⁹ https://bugzilla.redhat.com/show_bug.cgi?id=986090

³⁰ https://bugzilla.redhat.com/show_bug.cgi?id=882769

-
- BZ# [852293](#)³¹

Glusterd does not have a rollback mechanism to revert any actions that may have succeeded on some nodes and failed on the those that do not have the brick's parent dir. For example, setting the volume-id extended attribute may fail on some nodes (due to the absence of the directory for instance) 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. Either remove the brick directories or remove the glusterfs-related extended attributes.
2. Try creating the volume again.

- **Issues related to Gluster Samba**

- BZ# [994990](#)³²

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 writing operation because of the lock not being available.

Workaround: To avoid the issue, execute this command:

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

- **General issues**

- BZ# [986229](#)³³

In the Red Hat Storage sever, the ISCSI service begins by default.

- **yum update** command may fail when it is executed for the first time in Red Hat Virtual Storage Appliance 3.2 with the following error "GPG key retrieval failed: [Errno 14] Could not open/read file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary " .

Work Around: Run **yum update** command again and it works fine on subsequent runs.

- glusterfsd - Error return code is always 0 after daemonizing the process.

Due to this, scripts that mount glusterFS or start glusterFS process must not depend on its return value.

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

Contact Red Hat Support for more information on this issue.

- BZ# [865672](#)³⁴

³¹ https://bugzilla.redhat.com/show_bug.cgi?id=852293

³² https://bugzilla.redhat.com/show_bug.cgi?id=994990

³³ https://bugzilla.redhat.com/show_bug.cgi?id=986229

³⁴ https://bugzilla.redhat.com/show_bug.cgi?id=865672

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 on 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)³⁵

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 is not capable of differentiating between volumes deleted and volumes added in the absence of the other (conflicting) peers.

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

- BZ# [920002](https://bugzilla.redhat.com/show_bug.cgi?id=920002)³⁶

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

- BZ# [921263](https://bugzilla.redhat.com/show_bug.cgi?id=921263)³⁷

If glusterd is consuming lot of memory, restart the glusterd service.

- Due to enhancements in dynamic volume management (for scale-out feature), you may experience excessive memory usage in this release.

- BZ# [916834](https://bugzilla.redhat.com/show_bug.cgi?id=916834)³⁸

Quick-read translator returns stale file handles for certain pattern of the file access. When running `dbench` application on the mount point, `dbench: read fails on handle 10030` message is displayed.

Work Around: Execute `gluster volume set VOLNAME quick-read off` command to avoid this issue.

³⁵ https://bugzilla.redhat.com/show_bug.cgi?id=839213

³⁶ https://bugzilla.redhat.com/show_bug.cgi?id=920002

³⁷ https://bugzilla.redhat.com/show_bug.cgi?id=921263

³⁸ 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 2.0.

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 2.0 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 2.0 Console Administration Guide*.

4.2. Striped Volumes

Striped volumes stripes data across bricks in the volume. For best results, you should 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 Administration Guide.

4.3. Distributed-Striped Volumes

Distributed striped volumes stripe data across two or more nodes in the trusted storage pool. You should use distributed striped volumes where the requirement is to scale storage and in high concurrency environments where accessing very large files is critical.

For more information, refer to section *Creating Distributed Striped Volumes* of the Administration Guide.

4.4. Distributed-Striped-Replicated Volumes

Distributed striped replicated volumes distributes striped data across replicated bricks in the trusted storage pool. For best results, you should 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 section *Creating Distributed Striped Replicated Volumes* in the Administration Guide.

4.5. Striped-Replicated Volumes

Striped replicated volumes stripes data across replicated bricks in the trusted storage pool. For best results, you should 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 section *Creating Striped Replicated Volumes* in the Administration Guide.

4.6. Replicated Volumes with Replica Count greater than 2

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 volume with replica count more than 2 is under technology preview.

For more information, refer to section *Creating Replicated Volumes* in the 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 **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 section *Stopping Remove Brick Operation* in the Administration Guide

4.9. Read-only Volume

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

Appendix A. Revision History

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.

