

Package Information

- Location from which the packages are used
 - Unknown- lost to time, probably Gluster web site
- Package Info - version of glusterfs package installed

```
gluster_3.7]# gluster --version
```

```
glusterfs 3.7.11 built on Apr 27 2016 14:09:22
```

```
Repository revision: git://git.gluster.com/glusterfs.git
```

Copyright (c) 2006-2011 Gluster Inc. <<http://www.gluster.com>>**Cluster Information**

- Number of nodes in the cluster - 100
- Hostnames and IPs of the gluster Node [if it is not a security issue]
 - Hostname / IP will help developers in understanding & correlating with the logs
 - 10.1.1.0 network
 - eth1 Link encap:Ethernet HWaddr 2C:60:0C:8B:F0:59
 - inet addr:10.1.1.4 Bcast:10.255.255.255 Mask:255.0.0.0
 - inet6 addr: fe80::2e60:cff:fe8b:f059/64 Scope:Link
 - UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
 - RX packets:272515531346 errors:2 dropped:0 overruns:0 frame:1
 - TX packets:310738116338 errors:0 dropped:0 overruns:0 carrier:0
 - collisions:0 txqueuelen:1000
 - RX bytes:220224023147580 (200.2 TiB) TX bytes:386703576878881 (351.7 TiB)
 - Memory:c7700000-c771ffff
- Output of `gluster peer status`
 - `gluster peer status`
 - Number of Peers: 0
- Node IP, from which the "x" operation is done
 - "x" here means any operation that causes the issue **all**

Volume Information

- Number of volumes - 2
- Volume Names **data-volume & test-volume**
- Volume on which the particular issue is seen [if applicable] **data-volume**
- Type of volumes
- Volume options if available
- Output of `gluster volume info`
 - `[root@mseas-data2 ~]# gluster volume info`
 -
 - Volume Name: data-volume
 - Type: Distribute

- ID : 892d9e3a-b38c-4971-b96a-8e4a496685ba
- Status : completed
-
- Status of volume: test-volume
- Gluster process TCP Port RDMA Port Online Pid
- -----
- Brick mseas-data2:/home/gbrick_test_1 49156 0 Y 26594
- Brick mseas-data2:/home/gbrick_test_2 49157 0 Y 26613
- NFS Server on localhost N/A N/A N N/A
-
- Task Status of Volume test-volume
- -----
- There are no active volume tasks
-
- Get the statedump of the volume with the problem

```
$ gluster volume statedump <vol-name>
```

This dumps statedump per brick process in /var/run/gluster

File names:

mit_gluster_statedump.zip

NOTE: Collect statedumps from one gluster Node in a directory.

Repeat it in all Nodes containing the bricks of the volume. All the so collected directories could be archived, compressed and attached to bug

Brick Information

- xfs options when brick partition was done
 - This could be obtained with this command :

```
$ xfs_info /dev/mapper/vg1-brick
```

```
[root@mseas-data2 mapper]# xfs_info /mnt/brick1
```

```
meta-data=/dev/sda      isize=256  agcount=164, agsize=268435455 blks
```

```
        =                sectsz=4096  attr=2, projid32bit=0
```

```
data      =                bsize=4096  blocks=43949752320, imaxpct=1
```

```
        =                sunit=0    swidth=0 blks
```

```
naming   =version 2        bsize=4096  ascii-ci=0
```

```
log    =internal          bsize=4096 blocks=521728, version=2
      =                  sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none           extsz=4096 blocks=0, rtextents=0
```

```
[root@mseas-data2 mapper]# xfs_info /mnt/brick2
```

```
meta-data=/dev/sdb          isize=256  agcount=164, agsize=268435455 blks
      =                  sectsz=4096 attr=2, projid32bit=0
data    =                  bsize=4096 blocks=43949752320, imaxpct=1
      =                  sunit=0   swidth=0 blks
naming  =version 2         bsize=4096  ascii-ci=0
log     =internal          bsize=4096 blocks=521728, version=2
      =                  sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none           extsz=4096 blocks=0, rtextents=0
```

- Extended attributes on the bricks
 - This could be obtained with this command:

```
$ getfattr -d -m. -ehex /rhs/brick1/b1
```

No output

Client Information

- OS Type (Windows, RHEL)
- OS Version : In case of Linux distro get the following :

```
$ uname -r $ cat /etc/issue
```

uname -r

2.6.32-642.1.1.el6.x86_64

[root@mseas-data2 brick1]# cat /etc/issue

CentOS release 6.8 (Final)

Kernel \r on an \m

- Fuse or NFS Mount point on the client with output of mount commands
-
- Output of `df -Th` command

Tool Information

- If any tools are used for testing, provide the info/version about it
- if any IO is simulated using a script, provide the script

Logs Information

- You can check logs for check for issues/warnings/errors.
 - Self-heal logs
 - Rebalance logs
 - Glusterd logs
 - Brick logs
 - NFS logs (if applicable)
 - Samba logs (if applicable)
 - Client mount log
- Add the entire logs as attachment, if its very large to paste as a comment

SOS report for CentOS/Fedora

- Get the sosreport from the involved gluster Node and Client [in case of CentOS /Fedora]
 - In `/tmp/sos.jVfQOM`
- Add a meaningful name/IP to the sosreport, by renaming/adding hostname/ip to the sosreport name

The only problem with using gluster mounted via NFS is that it does not >> respect the group write permissions which we need.

One example, we have a common working directory dri_fleat in the gluster volume

```
drwxrwsr-x 22 root dri_fleat 4.0K May 1 15:14 dri_fleat
```

my user (phaley) does not own that directory but is a member of the group dri_fleat and should have write permissions. When I go to the nfs-mounted version and try to use the touch command I get the following

```
ibfdr-compute-0-4(dri_fleat)% touch dum  
touch: cannot touch `dum': Permission denied
```

One of the sub-directories under dri_fleat is "test" which phaley owns

```
drwxrwsr-x 2 phaley dri_fleat 4.0K May 1 15:16 test
```

Under this directory (mounted via nfs) user phaley can write

```
ibfdr-compute-0-4(test)% touch dum  
ibfdr-compute-0-4(test)%
```

I have put the packet captures in

<http://mseas.mit.edu/download/phaley/GlusterUsers/TestNFSmount/>