

## Issue Report

### OS distribution:

Fedora 11 / Linux

### Kernel Version:

2.6.29.4-167.fc11.x86\_64 #1 SMP Wed May 27 17:27:08 EDT 2009 x86\_64 x86\_64 x86\_64 GNU/Linux

### GlusterFS Version:

Repository revision: v2.0.7

### GlusterFS Setup:

I setup 5 machines (A,B,C,D,E machines), each one has 1x 1TB harddrive, Those 5 machines all has two NICs and connected to 2 subnet, one is for GlusterFS data transfer (glusterfs subnet - 10.106.105.\*) and another is for control the machines (control subnet). A,B,C,D's disk space are used by GlusterFS, except E. A,B,C,D,E all using the same configuration file. As you can see, the E's disk hasn't been used in the configuration. Only four copies, and used just one translator: replicate translator. And each machine can access to the mount point.

### GlusterFS Configuration:

```
# THIS IS THE SERVER-END CONFIGURATION
# Brick 1
volume posix
    type storage/posix
    option directory /mnt/disk1
end-volume

volume locks
    type features/locks
    subvolumes posix
end-volume

volume brick
    type performance/io-threads
    option thread-count 16
    subvolumes locks
end-volume

# Server
volume server
    type protocol/server
    option transport-type tcp/server
    option transport.socket.bind-address `ifconfig -a | grep "10.106.105." | awk '{print $2}' |
awk 'BEGIN {FS=":"};{print $2}`
    option transport.socket.listen-port 6996
    subvolumes brick
    option auth.addr.brick.allow *
end-volume

# SERVER-END CONFIGURATION ENDS

# THIS IS THE CLIENT-END CONFIGURATION
# 3 Disks Machines
# Machine 1
volume cbrick1
    type protocol/client
    option transport-type tcp
    option remote-host 10.106.105.150
    option remote-port 6996
    option ping-timeout 60
```

```

        option remote-subvolume brick
end-volume

# Machine 2
volume cbrick4
    type protocol/client
    option transport-type tcp
    option remote-host 10.106.105.151
    option remote-port 6996
    option ping-timeout 60
    option remote-subvolume brick
end-volume

# Machine 3
volume cbrick7
    type protocol/client
    option transport-type tcp
    option remote-host 10.106.105.152
    option remote-port 6996
    option ping-timeout 60
    option remote-subvolume brick
end-volume

# Machine 4
volume cbrick10
    type protocol/client
    option transport-type tcp
    option remote-host 10.106.105.153
    option remote-port 6996
    option ping-timeout 60
    option remote-subvolume brick
end-volume

# All the bricks declare finished

# Replicate part
volume repl
    type cluster/replicate
    subvolumes cbrick1 cbrick4 cbrick7 cbrick10
end-volume

# CLIENT END CONFIGURATION ENDS

```

## Test Descriptions:

I created 1000 files on the mount point (/home) when glusterfs on every machine is online, each file is 5MB.

My test has two separate procedures – simulate failure procedure and checking procedure, and both are running simultaneously.

Simulate failure procedure:

Step 1: Randomly choose 1 to 3 machines in A,B,C,D 4 machines, bring down their NIC that connected to Glusterfs subnet by using command `ifconfig ethX down` via ssh.

Step 2: Sleep for 60 seconds

Step 3: Bring up all the NICs that being bring down in Step 1. Then do `ls -laR <mount point>` in each machine that bring down the NIC.

Step 4: Goto Step 1

Checking procedure:

Two method, both methods are running on machine E, which purely acts like client:

ls method: Start the simulate failure procedure, execute command: "ls -laR <mount point> > /root/test\_result/testX", where X is number, this command is executed for 10000 times.

Md5 method: Before simulate failure procedure begins, generate a file list of all the files in mount point. Then start the failure procedure, and for every file in the file list do `stat <mount point>/<filename> > /root/md5\_test\_result/testX`, where X is number, this command is executed for 100 times

### Test result:

Apparently, 10000 times test, the size of output file of `ls -laR` suppose to be exactly the same, but unfortunately it's not. I found some outputs lack of some files (the file size is different), or seen some files name appears twice. Like this picture below, left is good test result, right is the bizarre output of `ls -laR`, missing files, even the file ".", duplicate file "file19", and loosen files:

Left Panel (Good Test Result)	Right Panel (Bizarre Output)
total 490416	total 478792
drwxr-xr-x 2 root root 4096 2009-10-30 11:48 .	drwxr-xr-x 12 root root 4096 2009-10-30 11:49 ..
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file0	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file0
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file1	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file1
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file10	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file10
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file11	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file11
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file12	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file12
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file13	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file12
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file14	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file14
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file15	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file15
	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file15
	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file16
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file16	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file16
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file17	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file17
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file18	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file17
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file19	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file18
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file2	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file19
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file20	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file19
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file21	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file21
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file22	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file22
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file23	
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file24	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file24
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file25	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file25
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file26	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file28
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file27	
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file28	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file28
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file29	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file29
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file3	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file3
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file30	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file30
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file31	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file31
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file32	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file32
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file33	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file33
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file34	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file34
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file35	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file35
	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file35
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file36	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file36
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file37	
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file38	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file38
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file39	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file39
	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file39
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file4	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file39
-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file40	-rw-r--r-- 1 root root 5004454 2009-10-30 11:47 file4

**Log file:**

See attached log files, 5 in all, those log files are generated with debug mode enabled