
Subject: Install GlusterFS filesystem in a VPS
Posted by [raskar](#) on Tue, 24 Feb 2009 01:29:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,
I'm trying to install and use GlusterFS on a VPS.
I succeed to install and run it on the physical host.
I install fuse without kernel support like tis tuto <http://wiki.openvz.org/FUSE>
In the other side fuse is correctly working on my physical host.
GlusterFS detect my fuse install and can compile and install.

My problem is when i'm trying to run the glusterFS server deamon, it give some strange errors :

```
2009-02-24 02:25:45 C [posix.c:3604:init] brick: Extended attribute not supported, exiting
2009-02-24 02:25:45 E [xlator.c:563:xlator_init_rec] xlator: initialization of volume 'brick' failed,
review your volfile again
2009-02-24 02:25:45 E [glusterfsd.c:465:_xlator_graph_init] glusterfs: initializing translator failed
2009-02-24 02:25:45 E [glusterfsd.c:1096:main] glusterfs: translator initialization failed. exiting
```

I used the same attributes to install on my physical host and it's works very well.
Somebody has ever succeed to install GlusterFS on a VPS ?

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [maratrus](#) on Tue, 24 Feb 2009 08:32:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

Quote:

```
2009-02-24 02:25:45 C [posix.c:3604:init] brick: Extended attribute not supported, exiting
```

Seems that indeed extended attributes are not supported inside VE but I think they are also not supported on the HN i.e. partition that is mounted to /vz. Could you possibly check this assumption?

```
tune2fs -l /dev/*
```

Filesystem features should contain ext_attr record.

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [raskar](#) on Tue, 24 Feb 2009 18:27:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

The /vz directory is on the same partition than my system (it's an integration server).
My partition is /dev/sda1 and I can run glusterFSd on my HN.

```
tune2fs 1.40-WIP (14-Nov-2006)
Filesystem volume name: <none>
Last mounted on: <not available>
Filesystem UUID: ddc5c511-e274-4c72-a9d5-5c34b7255e92
Filesystem magic number: 0xEF53
Filesystem revision #: 1 (dynamic)
[B]Filesystem features: has_journal ext_attr resize_inode dir_index filetype[B]
needs_recovery sparse_super large_file
Filesystem flags: signed directory hash
Default mount options: (none)
Filesystem state: clean
Errors behavior: Continue
Filesystem OS type: Linux
Inode count: 18792448
Block count: 37570002
Reserved block count: 1878500
Free blocks: 36303058
Free inodes: 18711149
First block: 0
Block size: 4096
Fragment size: 4096
Reserved GDT blocks: 1015
Blocks per group: 32768
Fragments per group: 32768
Inodes per group: 16384
Inode blocks per group: 512
Filesystem created: Mon Feb 9 11:37:38 2009
Last mount time: Mon Feb 9 20:11:30 2009
Last write time: Mon Feb 9 20:11:30 2009
Mount count: 5
Maximum mount count: 30
Last checked: Mon Feb 9 11:37:38 2009
Check interval: 15552000 (6 months)
Next check after: Sat Aug 8 12:37:38 2009
Reserved blocks uid: 0 (user root)
Reserved blocks gid: 0 (group root)
First inode: 11
Inode size: 128
Journal inode: 8
First orphan inode: 11733633
Default directory hash: tea
```

Directory Hash Seed: 245cffd2-af21-4e9d-928b-a969580c16a1
Journal backup: inode blocks

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [maratrus](#) on Wed, 25 Feb 2009 07:37:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

ok then are you able to strace that program?
http://wiki.openvz.org/Stracing_a_program

It would be useful to know how they're checking extended attribute support.

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [raskar](#) on Mon, 09 Mar 2009 23:38:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is what my strace look like
I've did the same on the physical server and I have the same output.(works on the physical server)

File Attachments

1) [gluster.txt](#), downloaded 581 times

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [maratrus](#) on Tue, 10 Mar 2009 08:14:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

According to your first message glusterfsd should write error messages but attached strace output doesn't contain anything related to writing messages.

Could you please clarify how did you issue the "strace" command?

Could you please try "strace -fF -o out.txt ...".

Anyway, could you also try
"vzctl set \$VEID --capability sys_admin:on"?

Subject: Re: Install GlusterFS filesystem in a VPS

Posted by [raskar](#) on Tue, 10 Mar 2009 17:39:20 GMT

[View Forum Message](#) <> [Reply to Message](#)

This is the arguments I used with strace:

```
strace -o gluster.out -ff glusterfsd -f glusterfs-srv.vol
```

I just did it again, and I have the same output file (I just renamed it).

This is the associated glusterFS log:

EDIT :

IT'S WORKS !!

The last command

```
vzctl set $VEID --capability sys_admin:on
```

What do this command ?

Subject: Re: Install GlusterFS filesystem in a VPS

Posted by [dr.m](#) on Sat, 20 Jun 2009 16:45:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi

I had the same problems and managed glusterfsd to work now. Thanks a lot!

But, I now try to mount the device inside a container

```
glusterfs --debug -f /etc/glusterfs/glusterfs.vol /storage
```

and get the following error(s):

```
fuse: failed to open /dev/fuse: Permission denied
[2009-06-20 18:36:29] D [fuse-bridge.c:2747:init] glusterfs-fuse: fuse_mount() failed with error
Permission denied on mount point /storage
```

ls -al /dev/fuse is

```
crw-rw-rw- 1 root fuse 10, 229 20. Jun 16:35 /dev/fuse
```

and ls -al /storage is

```
drwxr-xr-x  2 root root 4096 20. Jun 17:14 storage
```

the kernel of that HN has fuse loaded... a uname is

```
2.6.18-14-fza-amd64 #1 SMP Mon Jan 5 17:36:46 UTC 2009 x86_64 GNU/Linux
```

I'll attach the whole debug output.

Did you get the mount to work? How?
Any suggestions?

Thanks in advance.
Marcus

File Attachments

1) [glusterfs_debug.txt](#), downloaded 531 times

Subject: Re: Install GlusterFS filesystem in a VPS
Posted by [dr.m](#) on Mon, 22 Jun 2009 09:24:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi again,

found the solution on my own (<http://wiki.openvz.org/FUSE>)

I just needed to tune my container, by performing the following in my HN:

```
# vzctl set $VEID --devices c:10:229:rw --save  
# vzctl exec $VEID mknod /dev/fuse c 10 229
```
