

Glusterfs rpm

Centos GlusterFS



1. : <http://download.gluster.org/pub/gluster/glusterfs/3.3/3.3.1/CentOS/>
- 2.

```
$> yum install -y fuse
$> rpm -ivh glusterfs-3.3.1-1.el5.i386.rpm
$> rpm -ivh glusterfs-fuse-3.3.1-1.el5.i386.rpm
$> rpm -ivh glusterfs-devel-3.3.1-1.el5.i386.rpm
$> rpm -ivh glusterfs-server-3.3.1-1.el5.i386.rpm
```

3. gluster

```
/etc/rc.d/init.d/glusterd start
chkconfig glusterd on
\\*
/etc/glusterfs/, /var/lib/glusterd/
```

4. fuse

```
modprobe fuse
echo "modprobe fuse" >> /etc/rc.local
dmesg | grep -i fuse
```

fuse init (API version 7.10) <-- fuse

5. Storage Pool

1. 1. Management

```
gluster peer probe IP  
ef) gluster peer probe 192.168.150.19
```

6. Pool `IP`

```
[root@localhost]# gluster peer status  
Number of Peers: 1  
Hostname: 192.168.150.19  
Uuid: 0bb2e93b-1ab1-4420-91c0-e5bcf203586b  
State: Peer in Cluster (Connected)
```

7. Storage pool `IP`

```
gluster peer detach IP  
\\ef) gluster peer detach 192.168.150.19
```

8. Volume `IP`

1. Volume `IP` Distributed(`IP` , `IP`)
 - `IP` `IP` Brick `IP` `IP` `IP` `IP`
 - Distributed `IP` `IP` `IP` Brick `IP` `IP` , `IP` Brick `IP` `IP` `IP` `IP` `IP` `IP` .
2. replicated(`IP` , replication `IP` `IP` `IP`)
 - `IP` `IP` Brick `IP` `IP` `IP` `IP` `IP` , `IP` `IP` .
3. stripe(`IP`) `IP` `IP`
 1. `IP` `IP` `IP` `IP` Brick `IP` `IP` `IP`
 2. Stripe `IP` `IP` `IP` Brick `IP` `IP` `IP` , `IP` `IP` `IP`

9. volume `IP` (Distributed `IP`)

```
gluster volume create IP vmlnx001:/data vmlnx002:/data vmlnx003:/data vmlnx004:/data  
\\ef) gluster volume create data 192.168.150.18:/data 192.168.150.19:/data  
\\\\gluster> volume status data  
Status of volume: data  
Gluster process Port Online Pid  
\\-----  
\\Brick 192.168.150.18:/data 24011 Y 3245  
Brick 192.168.150.19:/data 24011 Y 2481  
NFS Server on localhost 38467 Y 3251  
NFS Server on 192.168.150.19 38467 Y 2487  
\\gluster> volume info data  
Volume Name: data
```

```
Type: Distribute
Volume ID: 556d6065-f888-4198-8782-65bc03979a0b
Status: Started
Number of Bricks: 3
Transport-type: tcp
Bricks:
Brick1: 192.168.150.18:/data
Brick2: 192.168.150.19:/data
Brick3: 192.168.150.16:/data
```

10. volume `replicate` (Replicate `replicate`)

```
gluster> volume create data replica 2 192.168.150.18:/data 192.168.150.19:/data
\\glusterfs volume data
gluster> volume info all
Volume Name: data
Type: Replicate
Volume ID: c2ecd1b8-708e-47d4-8f15-adcd1b081987
Status: Started
Number of Bricks: 1 x 2 = 2
Transport-type: tcp
Bricks:
Brick1: 192.168.150.19:/data
Brick2: 192.168.150.18:/data
```

11. volume `stripe` (stripe `stripe`)

```
gluster volume create data stripe 2 transport tcp 192.168.150.18:/data 192.168.150.19:/data
```

12. volume `replicate` `replicate` `replicate` volume `replicate` `replicate` , `replicate` `replicate` `replicate`

ef) gluster volume create data replica 3 192.168.150.18:/data 192.168.150.19:/data 192.168.150.16:/data

\\data or a prefix of it is already part of a volume

- `glusterfs` , `glusterfs` `glusterfs` `glusterfs` , `glusterfs` `glusterfs` `glusterfs` .

```
root@localhost data]# ls -al
drwxr-xr-x 4 root root 4096 12月 6 16:00 .
drwxr-xr-x 24 root root 4096 12月 6 16:01 ..
drw----- 7 root root 4096 12月 6 15:17 .glusterfs
```

```
drwx----- 2 root root 16384 12月 5 13:48 lost+found
```

```
ls -ld /data/.glusterfs
```

```
setfattr -x trusted.glusterfs.volume-id /data/
setfattr -x trusted.gfid /data/
rm -rf .glusterfs
```

13. 启动数据卷

```
gluster> volume start data
```

14. 设置数据卷的访问权限

```
gluster volume set <volume_name> auth.allow <IP地址>
ef) gluster> volume set data auth.allow 192.168.150.*
```

15. 停止并删除数据卷

```
gluster volume stop <volume_name>
gluster volume delete <volume_name>
\\ef)
gluster volume stop data
gluster volume delete data
```

16. 挂载数据卷

```
mount -t glusterfs <server_name>:<volume_name> <mount point>
mount -t glusterfs 192.168.150.18:/data /test/
\\[root@localhost]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2        4.9G  1.5G  3.2G  31% /
/dev/sda5        9.6G  151M  9.0G   2% /data
/dev/sda1        99M   12M   83M  13% /boot
tmpfs            1014M    0 1014M   0% /dev/shm
glusterfs#192.168.150.18:/data
                9.6G  151M  9.0G   2% /test
```

17. `192.168.150.18:/data` `/mnt` `glusterfs defaults,_netdev` `0 0`

1*. `192.168.150.18:/data` `/mnt` `glusterfs defaults,_netdev` `0 0` .

18. Volume `data`

`gluster> volume add-brick dist_vol 172.27.0.9:/data`

`gluster> volume add-brick data 192.168.150.16:/data`

`gluster> volume add-brick data 192.168.150.19:/data`

`gluster> volume status data`

`gluster> volume status data`

`gluster> volume status data`

Status of volume: data

Gluster process	Port	Online	Pid
Brick 192.168.150.18:/data	24011	Y	3245
Brick 192.168.150.19:/data	24011	Y	2481
NFS Server on localhost	38467	Y	3251
NFS Server on 192.168.150.16	38467	Y	4216
NFS Server on 192.168.150.19	38467	Y	2487

`gluster> volume status data`

`gluster> volume status data`

Status of volume: data

Gluster process	Port	Online	Pid
Brick 192.168.150.18:/data	24011	Y	3245
Brick 192.168.150.19:/data	24011	Y	2481
Brick 192.168.150.16:/data	24011	Y	4384
NFS Server on localhost	38467	Y	3268
NFS Server on 192.168.150.16	38467	Y	4390
NFS Server on 192.168.150.19	38467	Y	2501

19. `gluster> volume set data performance.write-behind-window-size 1024MB`

`gluster> volume set data performance.write-behind-window-size 1024MB`

`gluster> volume set data performance.cache-size 512MB`