

Introduction To Gluster

Thomas Cameron

RHCA, RHCSS, RHCDS, RHCVA, RHCX

Chief Architect, Central US
Red Hat

@thomsdcameron

thomas@redhat.com

Agenda

- What is Gluster?
 - Gluster Project
 - Red Hat and Gluster
 - What Does It Do?
 - How Does It Work?
- How Do I Install It?
 - Filesystems
 - Packages

Agenda

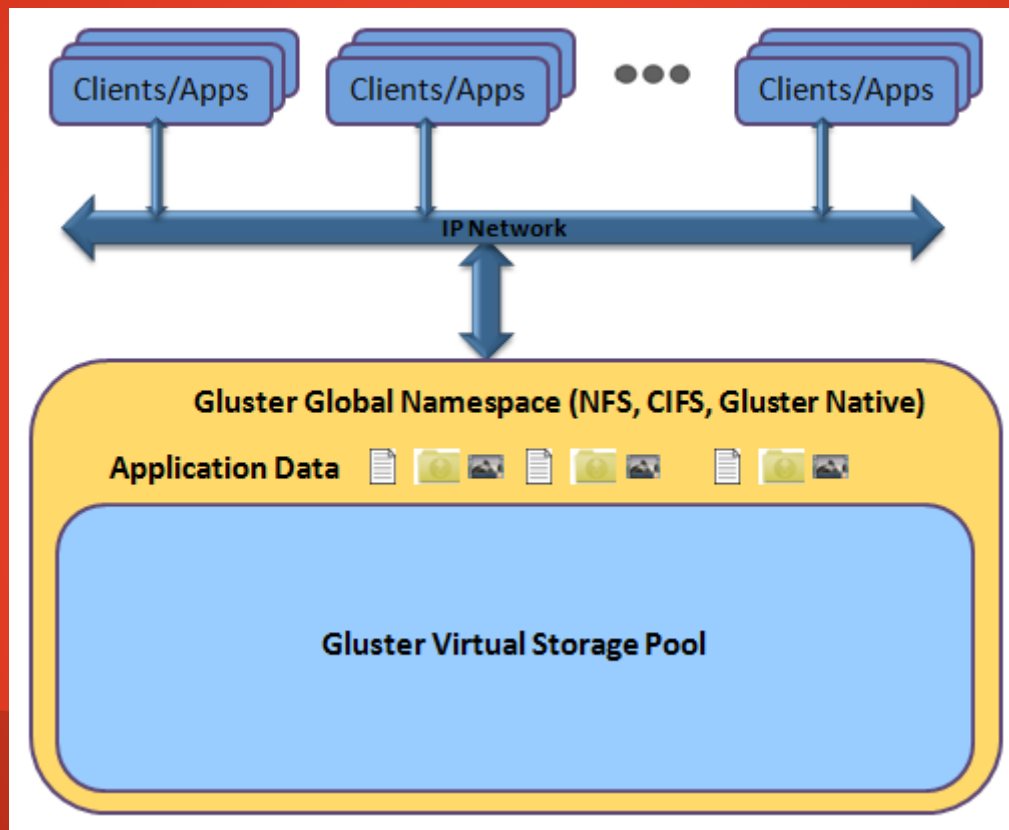
- How do I configure it?
 - Distributed
 - Replicated
- How Do I Connect to Gluster?
- Questions?

What is Gluster?

- Gluster Inc. was a software company that provided an open source platform for scale-out public and private cloud storage. The company was privately funded and headquartered in Sunnyvale, California with an engineering center in Bangalore, India. Gluster was funded by Nexus Venture Partners and Index Ventures. Gluster was acquired by Red Hat on October 7, 2011.
 - <http://en.wikipedia.org/wiki/Gluster>

What is Gluster?

- Red Hat bought Gluster for \$136,000,000 in 2011, and continues to make the software available as Open Source Software via the Gluster Project at <http://www.gluster.org/>
- GlusterFS is an open source, distributed file system capable of scaling to several petabytes (actually, 72 brontobytes!) and handling thousands of clients. GlusterFS clusters together storage building blocks over Infiniband RDMA or TCP/IP interconnect, aggregating disk and memory resources and managing data in a single global namespace. GlusterFS is based on a stackable user space design and can deliver exceptional performance for diverse workloads.



What is Gluster?

- GlusterFS supports standard clients running standard applications over any standard IP network. Figure 1, above, illustrates how users can access application data and files in a Global namespace using a variety of standard protocols.
- No longer are users locked into costly, monolithic, legacy storage platforms. GlusterFS gives users the ability to deploy scale-out, virtualized storage – scaling from terabytes to petabytes in a centrally managed and commoditized pool of storage.

What is Gluster?

- Attributes of GlusterFS include:
 - Scalability and Performance
 - High Availability
 - Global Namespace
 - Elastic Hash Algorithm
 - Elastic Volume Manager
 - Standards-based

What Does it Do?

- Gluster is software based network attached storage (NAS). It can serve Linux, Unix and Windows clients using the native Gluster client, NFS, or SMB, respectively.
- It can be used for any unstructured data - image files, log files, home directories, virtual machine images, object stores, etc.
- It is not for use with structured data like databases.

How Does It Work?

- Disks are added to bricks.
- Bricks are added to volumes.
- Volumes are exported via the native Gluster client, NFS, or SMB.

This Demo

- For this demo, I will use 4 VMs as servers, and my laptop as the client. Each of the server VMs has an 8GB root drive, and an 8GB drive which will be used as Gluster storage.
- Don't clone VMs with the Gluster software installed, it creates a UUID (/var/lib/glusterd/glusterd.info), and the nodes will not be able to see each other due to conflicts.

How Do I Install Gluster?

- Filesystem
 - Create a partition upon which you'll create the Gluster bricks
- Note that on production hardware, you'd probably want multiple spindles with RAID6 or RAID 10. For this quick start, one drive is fine.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# fdisk /dev/vdb
```

Welcome to fdisk (util-linux 2.25.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0x3abcc857.

Command (m for help): n

Partition type

- p primary (0 primary, 0 extended, 4 free)
- e extended (container for logical partitions)

Select (default p): p

Partition number (1-4, default 1):

First sector (2048-16383999, default 2048):

Last sector, +sectors or +size{K,M,G,T,P} (2048-16383999, default 16383999):

Created a new partition 1 of type 'Linux' and of size 7.8 GiB.

Command (m for help): w

The partition table has been altered.

Calling ioctl() to re-read partition table.

Syncing disks.

```
[root@gluster1 ~]#
```

```
root@gluster1:/root
File Edit View Search Terminal Help
Welcome to fdisk (util-linux 2.25.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0x3abcc857.

Command (m for help): n
Partition type
   p   primary (0 primary, 0 extended, 4 free)
   e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1):
First sector (2048-16383999, default 2048):
Last sector, +sectors or +size{K,M,G,T,P} (2048-16383999, default 16383999):

Created a new partition 1 of type 'Linux' and of size 7.8 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.

[root@gluster1 ~]#
```

```
root@gluster2:/root
File Edit View Search Terminal Help
Welcome to fdisk (util-linux 2.25.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0x482cbed7.

Command (m for help): n
Partition type
   p   primary (0 primary, 0 extended, 4 free)
   e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1):
First sector (2048-16383999, default 2048):
Last sector, +sectors or +size{K,M,G,T,P} (2048-16383999, default 16383999):

Created a new partition 1 of type 'Linux' and of size 7.8 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.

[root@gluster2 ~]#
```

```
root@gluster3:/root
File Edit View Search Terminal Help
Welcome to fdisk (util-linux 2.25.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0xbbe48fd4.

Command (m for help): n
Partition type
   p   primary (0 primary, 0 extended, 4 free)
   e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1):
First sector (2048-16383999, default 2048):
Last sector, +sectors or +size{K,M,G,T,P} (2048-16383999, default 16383999):

Created a new partition 1 of type 'Linux' and of size 7.8 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.

[root@gluster3 ~]#
```

```
root@gluster4:/root
File Edit View Search Terminal Help
Welcome to fdisk (util-linux 2.25.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0x659e0684.

Command (m for help): n
Partition type
   p   primary (0 primary, 0 extended, 4 free)
   e   extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-16383999, default 2048):
Last sector, +sectors or +size{K,M,G,T,P} (2048-16383999, default 16383999):

Created a new partition 1 of type 'Linux' and of size 7.8 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.

[root@gluster4 ~]#
```

How Do I Install Gluster?

- Filesystem
 - Use XFS as the underlying filesystem, with 512b inode size.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkfs.xfs -i size=512 /dev/vdb1
meta-data=/dev/vdb1          isize=512    agcount=4, agsize=511936 blks
                =                       sectsz=512   attr=2, projid32bit=1
                =                       crc=0      finobt=0
data        =                       bsize=4096  blocks=2047744, imaxpct=25
                =                       sunit=0     swidth=0 blks
naming      =version 2          bsize=4096  ascii-ci=0 ftype=0
log         =internal log     bsize=4096  blocks=2560, version=2
                =                       sectsz=512   sunit=0 blks, lazy-count=1
realtime    =none              extsz=4096  blocks=0, rtextents=0
[root@gluster1 ~]#
```



```
root@gluster1:/root
File Edit View Search Terminal Help
[root@gluster1 ~]# mkfs.xfs -i size=512 /dev/vdb1
meta-data=/dev/vdb1      isize=512    agcount=4, agsize=511936 blks
=                   sectsz=512   attr=2, projid32bit=1
=                   crc=0        finobt=0
data        =                   bsize=4096  blocks=2047744, imaxpct=25
=                   =                   sunit=0     swidth=0 blks
naming      =version 2          bsize=4096  ascii-ci=0  ftype=0
log         =internal log     bsize=4096  blocks=2560, version=2
=                   sectsz=512   sunit=0 blks, lazy-count=1
realtime    =none             extsz=4096  blocks=0, rtextents=0
[root@gluster1 ~]#
```

```
root@gluster2:/root
File Edit View Search Terminal Help
[root@gluster2 ~]# mkfs.xfs -i size=512 /dev/vdb1
meta-data=/dev/vdb1      isize=512    agcount=4, agsize=511936 blks
=                   sectsz=512   attr=2, projid32bit=1
=                   crc=0        finobt=0
data        =                   bsize=4096  blocks=2047744, imaxpct=25
=                   =                   sunit=0     swidth=0 blks
naming      =version 2          bsize=4096  ascii-ci=0  ftype=0
log         =internal log     bsize=4096  blocks=2560, version=2
=                   sectsz=512   sunit=0 blks, lazy-count=1
realtime    =none             extsz=4096  blocks=0, rtextents=0
[root@gluster2 ~]#
```

```
root@gluster3:/root
File Edit View Search Terminal Help
[root@gluster3 ~]# mkfs.xfs -i size=512 /dev/vdb1
meta-data=/dev/vdb1      isize=512    agcount=4, agsize=511936 blks
=                   sectsz=512   attr=2, projid32bit=1
=                   crc=0        finobt=0
data        =                   bsize=4096  blocks=2047744, imaxpct=25
=                   =                   sunit=0     swidth=0 blks
naming      =version 2          bsize=4096  ascii-ci=0  ftype=0
log         =internal log     bsize=4096  blocks=2560, version=2
=                   sectsz=512   sunit=0 blks, lazy-count=1
realtime    =none             extsz=4096  blocks=0, rtextents=0
[root@gluster3 ~]#
```

```
root@gluster4:/root
File Edit View Search Terminal Help
[root@gluster4 ~]# mkfs.xfs -i size=512 /dev/vdb1
meta-data=/dev/vdb1      isize=512    agcount=4, agsize=511936 blks
=                   sectsz=512   attr=2, projid32bit=1
=                   crc=0        finobt=0
data        =                   bsize=4096  blocks=2047744, imaxpct=25
=                   =                   sunit=0     swidth=0 blks
naming      =version 2          bsize=4096  ascii-ci=0  ftype=0
log         =internal log     bsize=4096  blocks=2560, version=2
=                   sectsz=512   sunit=0 blks, lazy-count=1
realtime    =none             extsz=4096  blocks=0, rtextents=0
[root@gluster4 ~]#
```

How Do I Install Gluster?

- Mount the new filesystem someplace intuitive
- Create a directory where we'll create the brick.
 - Don't create the brick at the root of the mount point. Use a subdirectory.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkdir -p /export/vdb1
```

```
[root@gluster1 ~]# ls /export/
```

```
vdb1
```

```
[root@gluster1 ~]# █
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkdir -p /export/vdb1
[root@gluster1 ~]# ls /export/
vdb1
[root@gluster1 ~]#
```

root@gluster2:/root

File Edit View Search Terminal Help

```
[root@gluster2 ~]# mkdir -p /export/vdb1
[root@gluster2 ~]# ls /export/
vdb1
[root@gluster2 ~]#
```

root@gluster3:/root

File Edit View Search Terminal Help

```
[root@gluster3 ~]# mkdir -p /export/vdb1
[root@gluster3 ~]# ls /export/
vdb1
[root@gluster3 ~]#
```

root@gluster4:/root

File Edit View Search Terminal Help

```
[root@gluster4 ~]# mkdir -p /export/vdb1
[root@gluster4 ~]# ls /export/
vdb1
[root@gluster4 ~]#
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkdir -p /export/vdb1
[root@gluster1 ~]# ls /export/
vdb1
[root@gluster1 ~]# mount /dev/vdb1 /export/vdb1/
[root@gluster1 ~]# mkdir /export/vdb1/brick
[root@gluster1 ~]# ls /export/vdb1/
brick
[root@gluster1 ~]# █
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkdir -p /export/vdb1
[root@gluster1 ~]# ls /export/
vdb1
[root@gluster1 ~]# mount /dev/vdb1 /export/vdb1/
[root@gluster1 ~]# mkdir /export/vdb1/brick
[root@gluster1 ~]# ls /export/vdb1/
brick
[root@gluster1 ~]#
```

root@gluster2:/root

File Edit View Search Terminal Help

```
[root@gluster2 ~]# mkdir -p /export/vdb1
[root@gluster2 ~]# ls /export/
vdb1
[root@gluster2 ~]# mount /dev/vdb1 /export/vdb1/
[root@gluster2 ~]# mkdir /export/vdb1/brick
[root@gluster2 ~]# ls /export/vdb1/
brick
[root@gluster2 ~]#
```

root@gluster3:/root

File Edit View Search Terminal Help

```
[root@gluster3 ~]# mkdir -p /export/vdb1
[root@gluster3 ~]# ls /export/
vdb1
[root@gluster3 ~]# mount /dev/vdb1 /export/vdb1/
[root@gluster3 ~]# mkdir /export/vdb1/brick
[root@gluster3 ~]# ls /export/vdb1/
brick
[root@gluster3 ~]#
```

root@gluster4:/root

File Edit View Search Terminal Help

```
[root@gluster4 ~]# mkdir -p /export/vdb1
[root@gluster4 ~]# ls /export/
vdb1
[root@gluster4 ~]# mount /dev/vdb1 /export/vdb1/
[root@gluster4 ~]# mkdir /export/vdb1/brick
[root@gluster4 ~]# ls /export/vdb1/
brick
[root@gluster4 ~]#
```

How Do I Install Gluster?

- Add the mount point to `/etc/fstab`

root@gluster1:/root

- □ ×

File Edit View Search Terminal Help

```
[root@gluster1 ~]# tail -1 /etc/mtab
/dev/vdb1 /export/vdb1 xfs rw,seclabel,relatime,attr2,inode64,noquota 0 0
[root@gluster1 ~]# tail -1 /etc/mtab >> /etc/fstab
[root@gluster1 ~]# █
```


root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# tail -1 /etc/mtab  
/dev/vdb1 /export/vdb1 xfs rw,seclabel,relatime,attr2,inode64,noquota 0 0  
[root@gluster1 ~]# tail -1 /etc/mtab >> /etc/fstab  
[root@gluster1 ~]#
```

root@gluster2:/root

File Edit View Search Terminal Help

```
[root@gluster2 ~]# tail -1 /etc/mtab  
/dev/vdb1 /export/vdb1 xfs rw,seclabel,relatime,attr2,inode64,noquota 0 0  
[root@gluster2 ~]# tail -1 /etc/mtab >> /etc/fstab  
[root@gluster2 ~]#
```

root@gluster3:/root

File Edit View Search Terminal Help

```
[root@gluster3 ~]# tail -1 /etc/mtab  
/dev/vdb1 /export/vdb1 xfs rw,seclabel,relatime,attr2,inode64,noquota 0 0  
[root@gluster3 ~]# tail -1 /etc/mtab >> /etc/fstab  
[root@gluster3 ~]#
```

root@gluster4:/root

File Edit View Search Terminal Help

```
[root@gluster4 ~]# tail -1 /etc/mtab  
/dev/vdb1 /export/vdb1 xfs rw,seclabel,relatime,attr2,inode64,noquota 0 0  
[root@gluster4 ~]# tail -1 /etc/mtab >> /etc/fstab  
[root@gluster4 ~]#
```

How Do I Install Gluster?

- Install the software

root@gluster1:/root



File Edit View Search Terminal Help

```
[root@gluster1 ~]# yum install glusterfs{,-server,-fuse,-geo-replication}
```

root@gluster1:/root

File Edit View Search Terminal Help

Running transaction check

Running transaction test

Transaction test succeeded

Running transaction (shutdown inhibited)

Installing : glusterfs-cli-3.5.3-1.fc21.x86_64 1/3

Installing : glusterfs-server-3.5.3-1.fc21.x86_64 2/3

Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.

Created symlink from /etc/systemd/system/multi-user.target.wants/glusterfsd.service to /usr/lib/systemd/system/glusterfsd.service.

Installing : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Verifying : glusterfs-server-3.5.3-1.fc21.x86_64 1/3

Verifying : glusterfs-cli-3.5.3-1.fc21.x86_64 2/3

Verifying : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Installed:

glusterfs-geo-replication.x86_64 0:3.5.3-1.fc21

glusterfs-server.x86_64 0:3.5.3-1.fc21

Dependency Installed:

glusterfs-cli.x86_64 0:3.5.3-1.fc21

Complete!

[root@gluster1 ~]#

```
root@gluster1:/root
File Edit View Search Terminal Help
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction (shutdown inhibited)
  Installing : glusterfs-cli-3.5.3-1.fc21.x86_64           1/3
  Installing : glusterfs-server-3.5.3-1.fc21.x86_64       2/3
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterfsd.service to /usr/lib/systemd/system/glusterfsd.service.
  Installing : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3
  Verifying  : glusterfs-server-3.5.3-1.fc21.x86_64       1/3
  Verifying  : glusterfs-cli-3.5.3-1.fc21.x86_64         2/3
  Verifying  : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Installed:
  glusterfs-geo-replication.x86_64 0:3.5.3-1.fc21
  glusterfs-server.x86_64 0:3.5.3-1.fc21

Dependency Installed:
  glusterfs-cli.x86_64 0:3.5.3-1.fc21

Complete!
[root@gluster1 ~]#
```

```
root@gluster2:/root
File Edit View Search Terminal Help
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction (shutdown inhibited)
  Installing : glusterfs-cli-3.5.3-1.fc21.x86_64           1/3
  Installing : glusterfs-server-3.5.3-1.fc21.x86_64       2/3
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterfsd.service to /usr/lib/systemd/system/glusterfsd.service.
  Installing : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3
  Verifying  : glusterfs-server-3.5.3-1.fc21.x86_64       1/3
  Verifying  : glusterfs-cli-3.5.3-1.fc21.x86_64         2/3
  Verifying  : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Installed:
  glusterfs-geo-replication.x86_64 0:3.5.3-1.fc21
  glusterfs-server.x86_64 0:3.5.3-1.fc21

Dependency Installed:
  glusterfs-cli.x86_64 0:3.5.3-1.fc21

Complete!
[root@gluster2 ~]#
```

```
root@gluster3:/root
File Edit View Search Terminal Help
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction (shutdown inhibited)
  Installing : glusterfs-cli-3.5.3-1.fc21.x86_64           1/3
  Installing : glusterfs-server-3.5.3-1.fc21.x86_64       2/3
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterfsd.service to /usr/lib/systemd/system/glusterfsd.service.
  Installing : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3
  Verifying  : glusterfs-server-3.5.3-1.fc21.x86_64       1/3
  Verifying  : glusterfs-cli-3.5.3-1.fc21.x86_64         2/3
  Verifying  : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Installed:
  glusterfs-geo-replication.x86_64 0:3.5.3-1.fc21
  glusterfs-server.x86_64 0:3.5.3-1.fc21

Dependency Installed:
  glusterfs-cli.x86_64 0:3.5.3-1.fc21

Complete!
[root@gluster3 ~]#
```

```
root@gluster4:/root
File Edit View Search Terminal Help
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction (shutdown inhibited)
  Installing : glusterfs-cli-3.5.3-1.fc21.x86_64           1/3
  Installing : glusterfs-server-3.5.3-1.fc21.x86_64       2/3
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterfsd.service to /usr/lib/systemd/system/glusterfsd.service.
  Installing : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3
  Verifying  : glusterfs-server-3.5.3-1.fc21.x86_64       1/3
  Verifying  : glusterfs-cli-3.5.3-1.fc21.x86_64         2/3
  Verifying  : glusterfs-geo-replication-3.5.3-1.fc21.x86_64 3/3

Installed:
  glusterfs-geo-replication.x86_64 0:3.5.3-1.fc21
  glusterfs-server.x86_64 0:3.5.3-1.fc21

Dependency Installed:
  glusterfs-cli.x86_64 0:3.5.3-1.fc21

Complete!
[root@gluster4 ~]#
```

How Do I Install Gluster?

- Turn on glusterd (elastic volume management daemon) and glusterfsd (Gluster server) and make them start at boot time.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# systemctl restart glusterd; systemctl restart glusterfsd
[root@gluster1 ~]# systemctl status glusterd; systemctl status glusterfsd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled)
   Active: active (running) since Thu 2015-02-19 15:06:23 PST; 11s ago
   Process: 2184 ExecStart=/usr/sbin/glusterd -p /run/glusterd.pid (code=exited,
status=0/SUCCESS)
   Main PID: 2185 (glusterd)
   CGroup: /system.slice/glusterd.service
           └─2185 /usr/sbin/glusterd -p /run/glusterd.pid
● glusterfsd.service - GlusterFS brick processes (stopping only)
   Loaded: loaded (/usr/lib/systemd/system/glusterfsd.service; enabled)
   Active: active (exited) since Thu 2015-02-19 15:06:23 PST; 11s ago
   Process: 2343 ExecStop=/bin/sh -c /bin/killall --wait glusterfsd || /bin/true
(code=exited, status=0/SUCCESS)
   Process: 2348 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
   Main PID: 2348 (code=exited, status=0/SUCCESS)
[root@gluster1 ~]#
```

```
root@gluster1:/root
File Edit View Search Terminal Help
[root@gluster1 ~]# systemctl restart glusterd; systemctl restart glusterfsd
[root@gluster1 ~]# systemctl status glusterd; systemctl status glusterfsd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled)
   Active: active (running) since Thu 2015-02-19 15:06:23 PST; 11s ago
   Process: 2184 ExecStart=/usr/sbin/glusterd -p /run/glusterd.pid (code=exited, status=0/SUCCESS)
   Main PID: 2185 (glusterd)
   CGroup: /system.slice/glusterd.service
           └─2185 /usr/sbin/glusterd -p /run/glusterd.pid
● glusterfsd.service - GlusterFS brick processes (stopping only)
   Loaded: loaded (/usr/lib/systemd/system/glusterfsd.service; enabled)
   Active: active (exited) since Thu 2015-02-19 15:06:23 PST; 11s ago
   Process: 2343 ExecStop=/bin/sh -c /bin/killall --wait glusterfsd || /bin/true (code=exited, status=0/SUCCESS)
   Process: 2348 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
   Main PID: 2348 (code=exited, status=0/SUCCESS)
[root@gluster1 ~]#
```

```
root@gluster2:/root
File Edit View Search Terminal Help
[root@gluster2 ~]# systemctl restart glusterd; systemctl restart glusterfsd
[root@gluster2 ~]# systemctl status glusterd; systemctl status glusterfsd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled)
   Active: active (running) since Thu 2015-02-19 15:06:52 PST; 4s ago
   Process: 1883 ExecStart=/usr/sbin/glusterd -p /run/glusterd.pid (code=exited, status=0/SUCCESS)
   Main PID: 1884 (glusterd)
   CGroup: /system.slice/glusterd.service
           └─1884 /usr/sbin/glusterd -p /run/glusterd.pid
● glusterfsd.service - GlusterFS brick processes (stopping only)
   Loaded: loaded (/usr/lib/systemd/system/glusterfsd.service; enabled)
   Active: active (exited) since Thu 2015-02-19 15:06:52 PST; 4s ago
   Process: 2041 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
   Main PID: 2041 (code=exited, status=0/SUCCESS)
[root@gluster2 ~]#
```

```
root@gluster3:/root
File Edit View Search Terminal Help
[root@gluster3 ~]# systemctl restart glusterd; systemctl restart glusterfsd
[root@gluster3 ~]# systemctl status glusterd; systemctl status glusterfsd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled)
   Active: active (running) since Thu 2015-02-19 15:06:53 PST; 3s ago
   Process: 1939 ExecStart=/usr/sbin/glusterd -p /run/glusterd.pid (code=exited, status=0/SUCCESS)
   Main PID: 1940 (glusterd)
   CGroup: /system.slice/glusterd.service
           └─1940 /usr/sbin/glusterd -p /run/glusterd.pid
● glusterfsd.service - GlusterFS brick processes (stopping only)
   Loaded: loaded (/usr/lib/systemd/system/glusterfsd.service; enabled)
   Active: active (exited) since Thu 2015-02-19 15:06:53 PST; 3s ago
   Process: 2097 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
   Main PID: 2097 (code=exited, status=0/SUCCESS)
[root@gluster3 ~]#
```

```
root@gluster4:/root
File Edit View Search Terminal Help
[root@gluster4 ~]# systemctl restart glusterd; systemctl restart glusterfsd
[root@gluster4 ~]# systemctl status glusterd; systemctl status glusterfsd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled)
   Active: active (running) since Thu 2015-02-19 15:06:54 PST; 3s ago
   Process: 1928 ExecStart=/usr/sbin/glusterd -p /run/glusterd.pid (code=exited, status=0/SUCCESS)
   Main PID: 1929 (glusterd)
   CGroup: /system.slice/glusterd.service
           └─1929 /usr/sbin/glusterd -p /run/glusterd.pid
● glusterfsd.service - GlusterFS brick processes (stopping only)
   Loaded: loaded (/usr/lib/systemd/system/glusterfsd.service; enabled)
   Active: active (exited) since Thu 2015-02-19 15:06:54 PST; 3s ago
   Process: 2086 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
   Main PID: 2086 (code=exited, status=0/SUCCESS)
[root@gluster4 ~]#
```


Tell The Nodes About Each Other

- `gluster peer probe [host]`
- `gluster peer probe [host]`
- `gluster peer probe [host]`
- ...

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster peer status
```

```
Number of Peers: 0
```

```
[root@gluster1 ~]# gluster peer probe gluster2.tc.redhat.com
```

```
peer probe: success.
```

```
[root@gluster1 ~]# gluster peer probe gluster3.tc.redhat.com
```

```
peer probe: success.
```

```
[root@gluster1 ~]# gluster peer probe gluster4.tc.redhat.com
```

```
peer probe: success.
```

```
[root@gluster1 ~]# gluster peer status
```

```
Number of Peers: 3
```

```
Hostname: gluster2.tc.redhat.com
```

```
Uuid: 52a53ddb-a5b5-4ff0-8651-db64c6292081
```

```
State: Peer in Cluster (Connected)
```

```
Hostname: gluster3.tc.redhat.com
```

```
Uuid: c67b6f90-092f-43ec-bb1e-00d98eb87878
```

```
State: Peer in Cluster (Connected)
```

```
Hostname: gluster4.tc.redhat.com
```

```
Uuid: 2a4ead03-7b68-4987-83b0-e6bad18fa2d5
```

```
State: Peer in Cluster (Connected)
```

```
[root@gluster1 ~]# █
```

```
root@gluster1:/root
File Edit View Search Terminal Help
[root@gluster1 ~]# gluster peer status
Number of Peers: 0
[root@gluster1 ~]# gluster peer probe gluster2.tc.redhat.com
peer probe: success.
[root@gluster1 ~]# gluster peer probe gluster3.tc.redhat.com
peer probe: success.
[root@gluster1 ~]# gluster peer probe gluster4.tc.redhat.com
peer probe: success.
[root@gluster1 ~]# gluster peer status
Number of Peers: 3

Hostname: gluster2.tc.redhat.com
Uuid: 52a53ddb-a5b5-4ff0-8651-db64c6292081
State: Peer in Cluster (Connected)

Hostname: gluster3.tc.redhat.com
Uuid: c67b6f90-092f-43ec-bb1e-00d98eb87878
State: Peer in Cluster (Connected)

Hostname: gluster4.tc.redhat.com
Uuid: 2a4ead03-7b68-4987-83b0-e6bad18fa2d5
State: Peer in Cluster (Connected)
[root@gluster1 ~]#
```

```
root@gluster2:/root
File Edit View Search Terminal Help
[root@gluster2 ~]# gluster peer status
Number of Peers: 0
[root@gluster2 ~]# gluster peer status
Number of Peers: 3

Hostname: 192.168.122.11
Uuid: f6afd58e-011d-42df-9553-e75735e44d48
State: Peer in Cluster (Connected)

Hostname: gluster3.tc.redhat.com
Uuid: c67b6f90-092f-43ec-bb1e-00d98eb87878
State: Peer in Cluster (Connected)

Hostname: gluster4.tc.redhat.com
Uuid: 2a4ead03-7b68-4987-83b0-e6bad18fa2d5
State: Peer in Cluster (Connected)
[root@gluster2 ~]#
```

```
root@gluster3:/root
File Edit View Search Terminal Help
[root@gluster3 ~]# gluster peer status
Number of Peers: 0
[root@gluster3 ~]# gluster peer status
Number of Peers: 3

Hostname: 192.168.122.11
Uuid: f6afd58e-011d-42df-9553-e75735e44d48
State: Peer in Cluster (Connected)

Hostname: gluster2.tc.redhat.com
Uuid: 52a53ddb-a5b5-4ff0-8651-db64c6292081
State: Peer in Cluster (Connected)

Hostname: gluster4.tc.redhat.com
Uuid: 2a4ead03-7b68-4987-83b0-e6bad18fa2d5
State: Peer in Cluster (Connected)
[root@gluster3 ~]#
```

```
root@gluster4:/root
File Edit View Search Terminal Help
[root@gluster4 ~]# gluster peer status
Number of Peers: 0
[root@gluster4 ~]# gluster peer status
Number of Peers: 3

Hostname: 192.168.122.11
Uuid: f6afd58e-011d-42df-9553-e75735e44d48
State: Peer in Cluster (Connected)

Hostname: gluster2.tc.redhat.com
Uuid: 52a53ddb-a5b5-4ff0-8651-db64c6292081
State: Peer in Cluster (Connected)

Hostname: gluster3.tc.redhat.com
Uuid: c67b6f90-092f-43ec-bb1e-00d98eb87878
State: Peer in Cluster (Connected)
[root@gluster4 ~]#
```

Create the Volume

- `gluster volume create [name] replica [count] transport [tcp|rdma] [host1:/path/to/brick] [host2:/path/to/brick]`
 - Note that in this first example, I leave out “replica” to create a purely distributed volume
- `gluster volume start [volume]`
- `gluster volume info [volume]` or `gluster volume info all`

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume info all
No volumes present
[root@gluster1 ~]# gluster volume create testvoll transport tcp gluster1:/export
/vdb1/brick/ gluster2:/export/vdb1/brick/
volume create: testvoll: success: please start the volume to access data
[root@gluster1 ~]# gluster volume start testvoll
volume start: testvoll: success
[root@gluster1 ~]# gluster volume info all

Volume Name: testvoll
Type: Distribute
Volume ID: 514b2ac7-774b-4333-aafd-9c1cfe68cbf5
Status: Started
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster1:/export/vdb1/brick
Brick2: gluster2:/export/vdb1/brick
[root@gluster1 ~]#
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume info all
No volumes present
[root@gluster1 ~]# gluster volume create testvol1 transport tcp gluster1:/export
/vdb1/brick/ gluster2:/export/vdb1/brick/
volume create: testvol1: success: please start the volume to access data
[root@gluster1 ~]# gluster volume start testvol1
volume start: testvol1: success
[root@gluster1 ~]# gluster volume info all
```

```
Volume Name: testvol1
Type: Distribute
Volume ID: 514b2ac7-774b-4333-aafd-9c1cfe68cbf5
Status: Started
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster1:/export/vdb1/brick
Brick2: gluster2:/export/vdb1/brick
[root@gluster1 ~]#
```

root@gluster2:/root

File Edit View Search Terminal Help

```
[root@gluster2 ~]# gluster volume info all

Volume Name: testvol1
Type: Distribute
Volume ID: 514b2ac7-774b-4333-aafd-9c1cfe68cbf5
Status: Started
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster1:/export/vdb1/brick
Brick2: gluster2:/export/vdb1/brick
[root@gluster2 ~]#
```

root@gluster3:/root

File Edit View Search Terminal Help

```
[root@gluster3 ~]# gluster volume info all
```

```
Volume Name: testvol1
Type: Distribute
Volume ID: 514b2ac7-774b-4333-aafd-9c1cfe68cbf5
Status: Started
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster1:/export/vdb1/brick
Brick2: gluster2:/export/vdb1/brick
[root@gluster3 ~]#
```

root@gluster4:/root

File Edit View Search Terminal Help


```
[root@gluster4 ~]# gluster volume info all
```

```
Volume Name: testvol1
Type: Distribute
Volume ID: 514b2ac7-774b-4333-aafd-9c1cfe68cbf5
Status: Started
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster1:/export/vdb1/brick
Brick2: gluster2:/export/vdb1/brick
[root@gluster4 ~]#
```

Configuring Gluster

- Note that there is nothing in the bricks

root@gluster1:/root



File Edit View Search Terminal Help

```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
[root@gluster1 ~]# █
```


Using Gluster

- From a client, mount the gluster export. In my case, I used my laptop
 - `mount -t glusterfs [server:volume] /local/path`
- Mount the export, then write some files. In this case, 10 files.

root@t540p:/root

File Edit View Search Terminal Help

```
[root@t540p ~]# mount -t glusterfs gluster1:testvoll /mnt/gluster/
```

```
[root@t540p ~]# df -h /mnt/gluster/
```

Filesystem	Size	Used	Avail	Use%	Mounted on
gluster1:testvoll	16G	65M	16G	1%	/mnt/gluster

```
[root@t540p ~]# for i in $(seq 1 10); do echo $i > /mnt/gluster/$i; done
```

```
[root@t540p ~]# ls /mnt/gluster/
```

```
1 10 2 3 4 5 6 7 8 9
```

```
[root@t540p ~]# █
```

Using Gluster

- On the servers, you'll see about $1/N$ of the files you just created in the brick directory, where N is the number of replicas you have. In this case, two.

root@gluster1:/root

File Edit View Search Terminal Help


```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
1 5 7 8 9
```

```
[root@gluster1 ~]#
```

root@gluster2:/root



File Edit View Search Terminal Help

```
[root@gluster2 ~]# ls /export/vdb1/brick/
```

```
10 2 3 4 6
```

```
[root@gluster2 ~]#
```

Using Gluster

- Now I'll remove this volume and create a replicated one.

root@gluster1:/root

- □ ×

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume stop testvol1
Stopping volume will make its data inaccessible. Do you want to continue? (y/n)
y
volume stop: testvol1: success
[root@gluster1 ~]# █
```

root@gluster1:/root



File Edit View Search Terminal Help


```
[root@gluster1 ~]# gluster volume delete testvol1
```

```
Deleting volume will erase all information about the volume. Do you want to continue? (y/n) y
```

```
volume delete: testvol1: success
```

```
[root@gluster1 ~]#
```


root@gluster1:/root




File Edit View Search Terminal Help

```
[root@gluster1 ~]# rm -rf /export/vdb1/brick/
```

```
[root@gluster1 ~]# █
```

root@gluster1:/root



File Edit View Search Terminal Help

```
[root@gluster1 ~]# mkdir /export/vdb1/brick/
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume create testvol2 transport tcp replica 2 gluster1:/export/vdb1/brick/ gluster2:/export/vdb1/brick/  
volume create: testvol2: success: please start the volume to access data  
[root@gluster1 ~]# gluster volume start testvol2  
volume start: testvol2: success  
[root@gluster1 ~]#
```

Using Gluster

- Again, mount the volume from a client.
- Note that this time, only 8GB are available (replicated)

root@t540p:/root

File Edit View Search Terminal Help

```
[root@t540p ~]# mount -t glusterfs gluster1:testvol2 /mnt/gluster/
```

```
[root@t540p ~]# df -h /mnt/gluster/
```

Filesystem	Size	Used	Avail	Use%	Mounted on
gluster1:testvol2	7.9G	33M	7.8G	1%	/mnt/gluster

```
[root@t540p ~]#
```

root@t540p:/root

- □ ×

File Edit View Search Terminal Help

```
[root@t540p ~]# for i in $(seq 1 10); do echo $i > /mnt/gluster/$i; done
```

```
[root@t540p ~]# ls /mnt/gluster/
```

```
1 10 2 3 4 5 6 7 8 9
```

```
[root@t540p ~]# █
```

Using Gluster

- This time, all of the files exist in each brick on all replicas.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
1 10 2 3 4 5 6 7 8 9
```

```
[root@gluster1 ~]#
```


root@gluster2:/root

_ □ ×

File Edit View Search Terminal Help

```
[root@gluster2 ~]# ls /export/vdb1/brick/
```

```
1 10 2 3 4 5 6 7 8 9
```

```
[root@gluster2 ~]#
```

Using Gluster

- I'll add more replicas to the volume with the gluster volume add-brick command.

root@gluster1:/root

- □ ×

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume add-brick testvol2 replica 4 gluster3:/export/  
vdb1/brick/ gluster4:/export/vdb1/brick/  
volume add-brick: success  
[root@gluster1 ~]# █
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# gluster volume info
```

Volume Name: testvol2

Type: Replicate

Volume ID: 4742df9f-bfe7-4eea-a3a2-b2e2cddb47b4

Status: Started

Number of Bricks: 1 x 4 = 4

Transport-type: tcp

Bricks:

Brick1: gluster1:/export/vdb1/brick

Brick2: gluster2:/export/vdb1/brick

Brick3: gluster3:/export/vdb1/brick

Brick4: gluster4:/export/vdb1/brick

```
[root@gluster1 ~]#
```

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
1 10 2 3 4 5 6 7 8 9
```

```
[root@gluster1 ~]#
```

Using Gluster

- Note that until the first write happens, the new bricks are empty.

root@gluster3:/root



File Edit View Search Terminal Help

```
[root@gluster3 ~]# ls /export/vdb1/brick/
```

```
[root@gluster3 ~]#
```

Using Gluster

- We write from the client (files 11 - 20)

root@t540p:/root

- □ ×

File Edit View Search Terminal Help

```
[root@t540p ~]# for i in $(seq 11 20); do echo $i > /mnt/gluster/$i; done
```

```
[root@t540p ~]# ls /mnt/gluster/
```

```
1 10 11 12 13 14 15 16 17 18 19 2 20 3 4 5 6 7 8 9
```

```
[root@t540p ~]#
```

root@gluster1:/root



File Edit View Search Terminal Help

```
[root@gluster1 ~]# ls /export/vdb1/brick/
```

```
1 10 11 12 13 14 15 16 17 18 19 2 20 3 4 5 6 7 8 9
```

```
[root@gluster1 ~]#
```

root@gluster3:/root

File Edit View Search Terminal Help

```
[root@gluster3 ~]# ls /export/vdb1/brick/
```

```
[root@gluster3 ~]# ls /export/vdb1/brick/
```

```
1 10 11 12 13 14 15 16 17 18 19 2 20 3 4 5 6 7 8 9
```

```
[root@gluster3 ~]#
```

How Do I Connect to Gluster?

- You've seen how to connect using the native Linux Gluster client:
 - `mount -t glusterfs`
- For NFS, you have to export the Gluster share via NFS from one of the nodes.

root@gluster1:/root

File Edit View Search Terminal Help

```
[root@gluster1 ~]# cat /etc/exports
/mnt/glusternfs *(fsid=0,rw)
[root@gluster1 ~]# systemctl enable rpcbind
[root@gluster1 ~]# systemctl enable nfs-server
[root@gluster1 ~]# systemctl start rpcbind
[root@gluster1 ~]# systemctl start nfs-server
[root@gluster1 ~]# showmount -e
Export list for gluster1.tc.redhat.com:
/mnt/glusternfs *
[root@gluster1 ~]# █
```

root@t540p:/root

File Edit View Search Terminal Help

```
[root@t540p ~]# showmount -e gluster1
```

```
Export list for gluster1:
```

```
/mnt/glusternfs *
```

```
[root@t540p ~]# umount /mnt/gluster
```

```
[root@t540p ~]# mount -t nfs gluster1:/mnt/glusternfs /mnt/gluster/
```

```
[root@t540p ~]# ls /mnt/gluster/
```

```
1 10 11 12 13 14 15 16 17 18 19 2 20 3 4 5 6 7 8 9
```

```
[root@t540p ~]# █
```

How Do I Connect to Gluster?

- For SMB, set up Samba in the same fashion - local Gluster mount exported via Samba (out of scope for today due to time constraints).

Questions?

Thank You!

