

Gluster FS

Ling Ho
SLAC National Accelerator Laboratory
January 12, 2012

Linux File/Storage Services

File System	Ext3, Ext4, XFS, ReiserFS, etc.
LVM	Disk consolidations / stripings
NFS	Network Sharings
RAID 1	High Availability
<i>DRBD</i>	Mirrorings, Backups
Quota	Limit Usage based on Size
ACL	Limit Usage based on User

What GlusterFS gives

File System	Works as a layer on top of any FS
LVM	Consolidates disks on 1 or multiple Linux nodes. Stripe across nodes.
NFS	Supports NFS, CIFS, and Native Client. Single or Multiple mount objects.
RAID 1	Replication (Synchronous)
<i>DRBD</i>	Geo-Replication (Asynchronous)
Quota	Control Directory Size
ACL	Posix Compliance

Gluster Server

5 RPMS to install

All command line driven, no file to edit

```
create file system and mount (mkfs /dev/sdb; mount /dev/sdb  
/brick1)  
gluster volume create vol1 transport tcp, rdma server1:  
/brick1  
gluster volume add-brick vol1 server1:/brick2  
gluster volume add-brick vol1 server2:/brick1  
...  
gluster start volume vol1
```

Gluster Client

Mount via NFS, CIFS or Native Client

~5 RPMS to install

All command line driven, no file to edit

NFS:

```
mount server1:/vol1 /data
```

Native Client:

```
mount -t glusterfs -o acl,transport=tcp server1:vol1  
/data
```

Performance

Tests results by Igor Gaponenko

READ

1 file, 1 server, 1 client node via Infiniband	726MB/s
1 file, 1 server, 20 client nodes via Infiniband	3700MB/s
80 files, 5 servers, 20 client nodes via Infiniband	9313/MBs

WRITE

1 file, 1 server, 1 client node via Infiniband	927MB/s
80 files, 5 servers, 20 client nodes via Infiniband	7439/MBs

Features

Works with any file system, even over LVM

All bricks do not have to be the same

Files can be accessed on the servers

Comes with rebalance, migration tools

Comes with server access control

Pretty good NFS server

Linux File/Storage Services

File System	Ext3, Ext4, XFS, ReiserFS, etc.
LVM	Disks consolidations
NFS	Network Sharings
RAID 1	High Availability
<i>DRBD</i>	Mirrorinsg, Backups
Quota	Limit Usage based on Size
ACL	Limit Usage based on User
GLUSTERFS	All of the Above