An Introduction To Gluster

Ryan Matteson
matty91@gmail.com
http://prefetch.net

Presentation Overview

- Tonight I am going to give an overview of Gluster, and how you can use it to create scalable, distributed file systems
- I love interactive presentations, so please ask if you aren't sure of something!

What Is Gluster?

- Gluster is an open source, scalable, distributed cluster file system capable of scaling to several brontobytes and thousands of clients
- Typically combined with commodity servers and storage to form massive storage networks

Gluster Features

- Gluster has a number of features that favorably tip the geek scale:
  - Global namespace
  - Clustered storage
  - Modular and stackable
  - Highly available storage
  - Built in replication and geo-replication
  - Self-healing
  - The ability to re-balance data

Gluster Terminology

- Four main concepts:
  - Bricks - storage units which consist of a server and directory path (i.e., server:/export)
  - Translators - modules that are chained together to move data from point a to point b
  - Trusted Storage Pool – a trusted network of servers that will host storage resources
  - Volumes - collection of bricks with a common redundancy requirement

Putting Things Together

- Trusted storage pools contain one or more storage servers that will host Gluster volumes
- A brick contains the name of a trusted storage server and a directory on the server where data will be read and written by clients
- Bricks are combined into volumes based on performance and reliability requirements
- Volumes are shared with Gluster clients through CIFS, NFS or the Gluster file system

Gluster Volume Types

- Gluster supports a number of volumes types, each providing different availability and performance characteristics:
  - ***Distributed*** – Files are distributed across bricks in the cluster
  - ***Replicated*** – Files are replicated across one or more bricks in the cluster
  - ***Striped*** – Stripes data across one or more bricks
  - ***Distributed replicated*** – Distributes files across replicated bricks in a cluster
  - ***Distributed striped*** – Stripes data across two or more nodes in the cluster

Which Volume Type Should I Use?

- From the official Gluster documentation:
  - Use distributed volumes where the requirement is to scale storage and the redundancy is either not important or is provided by other hardware/software layers
  - Use replicated volumes in environments where high-availability and high-reliability are critical
  - Use striped volumes only in high concurrency environments accessing very large files
  - Use distributed striped volumes where the requirement is to scale storage and in high concurrency environments accessing very large files
  - Use distributed replicated volumes in environments where the requirement is to scale storage and high-reliability is critical. Distributed replicated volumes offer improved read performance in most environments

Getting Gluster Working

- Seven step process:
  - Install the Gluster packages
  - Start the Gluster services
  - Create a trusted storage pool
  - Create new volumes
  - Start volumes
  - Lock down who can see the volumes
  - Mount the volumes on clients

Installing Gluster

- Three methods available:
  - For Fedora 16+ you can install via yum:

    ***$ yum install glusterfs flusterfs-fuse ***
    glusterfs-server glusterfs-vim glusterfs-devel
  - For RPM distributions rpmbuild is your friend:

    $ rpmbuild -ta glusterfs-version.tar.gz
  - Configure, make and make install also works
  - You can run `gluster –V` to verify your installation is complete and functional

Enabling Gluster Services

- The glusterd service needs to be started prior to using Gluster
- Starting Gluster on RHEL, CentOS and Fedora is crazy easy:

```
$ chkconfig glusterd on
$ service glusterd start
```

Adding Storage Servers To A Trusted Storage Pool

- A trusted storage pool consists of one or more servers, and each server can contain one or more bricks
- To add a server to a trusted storage pool you can run `gluster peer probe` followed by the hostname or IP of the server to add:

  $ gluster peer probe gluster02.prefetch.net
  Probe successful
- You can view cluster status with `gluster peer status`:

  $ gluster peer status
  Number of Peers: 1

  Hostname: gluster02.prefetch.net
  Uuid: 8667f377-5736-431b-b905-b607873035f0
  State: Peer in Cluster (Connected)

Creating Volumes

- You can create a volume with `gluster volume create <options>`:

  ***$ gluster volume create glustervol01 ***
  ***replica 2 transport tcp ***
  ***gluster01:/gluster/vol01 ***
  gluster02:/gluster/vol01

- In the example above I created a replicated volume named glustervol01, it contains two bricks and has a replica value of 2 to tell Gluster I want my data mirrored to two bricks

Starting Gluster Volumes

- Volumes need to be started after creation:

  $ gluster volume start glustervol01
- You can run `gluster volume info` to view volume status:

  $ gluster volume info
  Volume Name: glustervol01
  Type: Replicate
  Status: Created
  Number of Bricks: 2
  Transport-type: tcp
  Bricks:
  Brick1: gluster01.prefetch.net:/gluster/vol01
  Brick2: gluster02.prefetch.net:/gluster/vol01
  Options Reconfigured:
  auth.allow: 192.168.1.*

Mounting A Gluster File System

- You can mount a Gluster file system on a client with the mount command:

  $ mount -t glusterfs fedora-cluster01:/glustervol01 /gluster
- The server used in the mount command is only used to retrieve information about the Gluster volume
- Once mounted the client will interact with all of the bricks based on he volume type

Securing Gluster

- This is currently a major wart in the current stable release of Gluster
- Clients are authenticated based on IP address ranges, which we all know is less than ideal ☹
- Work is actively underway to:
  - Add certificate based authentication
  - Introduce an encryption translator

Securing Gluster (cont.)

- But all is not lost, we can still take a couple of actions to improve security:
  - Separate Gluster traffic from your production network traffic
  - Utilize iptables to limit who can talk to your Gluster trusted storage servers
  - Configure Gluster to only allow mounts from specific clients or networks:

```
$ gluster volume set glustervol01 \
  auth.allow 192.168.1.*
  Set volume successful
```

Conclusion

- Gluster provides an amazing amount of coolness for a relatively new file system
- Gluster is still in its infancy, and there are some growing pains. Once these are addressed Gluster will truly be amazing!
- It’s free, open source, so go grab a copy and start playing with it. You’ll love it!

Questions?

References

- Official Gluster website:

  http://gluster.org
- Posts I've written about Gluster:

  http://prefetch.net/blog/index.php/category/linux-gluster/
- The future of Gluster:

  http://redhatstorage.redhat.com/2012/01/27/the-future-of-glusterfs-slides/