

Step-by-Step Guide: Setup GlusterFS

If you dont know what GlusterFS is what what its for, you may consider check out this [Post](#).

Step-by-Step Guide to Setting Up GlusterFS on a 3-Node Cluster

This guide will walk you through setting up GlusterFS on a 3-node cluster using Raspberry Pis, with the IP addresses: `192.168.0.10`, `192.168.0.11`, and `192.168.0.12`. GlusterFS will be used to create a distributed, replicated file system across these nodes. Follow these steps to get your cluster up and running.

Step 1: Install GlusterFS on All Nodes

First, you need to install the required packages on **all nodes**. The `software-properties-common` package allows managing repositories, and `glusterfs-server` installs the GlusterFS server.

```
sudo apt install software-properties-common glusterfs-server -y
```

This ensures that every node in your cluster has GlusterFS installed and ready to share and replicate files.

Step 2: Enable Automatic Start of GlusterFS on Reboot

To ensure GlusterFS starts automatically after each reboot, you need to enable the `glusterd` service (GlusterFS daemon) on all nodes. Run this command on each node:

```
sudo systemctl start glusterd && sudo systemctl enable glusterd
```

This command starts the Gluster service immediately and ensures it starts automatically on system reboot.

Step 3: Peer Probe the Nodes

Now, log in to the main node (`192.168.0.10`) via SSH as root. You need to "peer probe" the other nodes to add them to the GlusterFS cluster. The peer probe is a command that connects the other nodes to the Gluster network, allowing them to participate in the distributed file system.

On the main node (`192.168.0.10`), run the following commands:

```
gluster peer probe 192.168.0.11
gluster peer probe 192.168.0.12
```

The `peer probe` command tells the main node to reach out to the other nodes, establishing a connection and syncing them into the cluster.

Step 4: Create the GlusterFS Volume

A volume in GlusterFS is essentially a storage pool made up of directories on different nodes. Here, we'll create a **replicated** volume across all three nodes. Replication ensures that the same data is stored on all nodes, making it resilient to failures.

Run the following command on the main node (`192.168.0.10`):

```
gluster volume create [glustertype] replica 3 192.168.0.10:/root/gluster 192.168.0.11:/root/gluster
192.168.0.12:/root/gluster force
```

Explanation:

- `[glustertype]` is the name you want to give to your GlusterFS volume.
- `replica 3` specifies that this is a replicated volume across all 3 nodes.
- The paths after each node's IP specify where the volume will be stored on each node (`/root/gluster`).

The `force` option is used to bypass potential warnings (such as creating volumes in a root directory).

Step 5: Start the GlusterFS Volume

Once the volume is created, you need to start it. This also needs to be done on the main node:

```
gluster volume start [glustername]
```

Replace `[glustername]` with the name of the volume you created in the previous step.

Step 6: Create a Mount Directory on Each Node

On **each node**, create a directory where you will mount the GlusterFS volume. For this guide, we'll use `/mnt/glustermount` as the mount point.

Run this command on all nodes:

```
sudo mkdir -p /mnt/glustermount
```

This ensures that the GlusterFS volume has a place to be mounted on each node.

Step 6.1: Mount the GlusterFS Volume

Now, mount the GlusterFS volume on the `/mnt/glustermount` folder. Run this command on each node:

```
sudo mount.glusterfs localhost:[glustername] /mnt/glustermount
```

Replace `[glustername]` with the name of your volume. This mounts the volume, making it accessible from the `/mnt/glustermount` directory on each node.

To Setup Automatic Mount of GlusterFS on Boot check out this [Post](#),

Your 3-node GlusterFS cluster should now be up and running, providing a replicated and distributed file system that's resilient and ready for use!

Revision #15

Created 8 September 2024 12:46:35 by aeoneros

Updated 9 September 2024 09:21:02 by aeoneros