

# Step-by-Step Guide: How to Mount GlusterFS on Boot

To ensure that your GlusterFS volume is automatically mounted at boot, you'll need to make some adjustments. By default, GlusterFS volumes might not mount properly on startup due to timing issues, where the Gluster service (`glusterd`) hasn't fully started yet. The following steps provide a workaround that ensures your GlusterFS volume is mounted after boot.

This guide is based on insights from [this ServerFault discussion](#), modified to suit your 3-node Raspberry Pi setup.

## Step 1: Modify the GlusterFS Service to Wait Before Mounting

First, you need to add a script that ensures the GlusterFS service has fully started before the system attempts to mount the volume.

Run the following command to edit the `glusterd.service`:

```
sudo systemctl edit glusterd.service
```

This will open a text editor. Add the following configuration to ensure a delay in mounting the GlusterFS volume:

```
[Service]
ExecStartPost=/usr/local/sbin/glusterfs-wait
```

Save and exit the editor.

## Step 2: Create the `glusterfs-wait` Script

Next, create a script that checks if the GlusterFS volume is ready to be mounted. This script will be called after the Gluster service starts and will retry the mount command until it succeeds.

Create the script file:

```
sudo nano /usr/local/sbin/glusterfs-wait
```

Paste the following content into the file:

```
#!/bin/bash
FAIL=1
until [ $FAIL -eq 0 ]; do
    gluster volume status all
    FAIL=$?
    test $FAIL -ne 0 && sleep 1
done
exit 0
```

Save and close the file.

This script will keep checking the status of the GlusterFS volume every second until the volume is ready.

## Step 3: Make the Script Executable

You need to make the script executable so it can run at startup.

Run this command:

```
sudo chmod +x /usr/local/sbin/glusterfs-wait
```

## Step 4: Create a Mount Override

To ensure that the GlusterFS mount happens after the `glusterd.service` is up and running, you need to create a mount override.

For the example folder `/mnt/glustermount`, run the following command:

```
sudo systemctl edit mnt-glustermount.mount
```

In the text editor, paste the following content:

```
[Unit]
```

```
After=glusterd.service
```

```
Wants=glusterd.service
```

This ensures that the mount service will wait for the `glusterd` service to start.  
Save and close the editor.

## Step 5: Reload the Systemd Daemon and Reboot

To apply the changes, reload the systemd daemon with the following command:

```
sudo systemctl daemon-reload
```

Then, reboot your system:

```
sudo reboot
```

## Step 6: Verify GlusterFS Mount After Reboot

After the system reboots, check the status of the GlusterFS service and ensure the volume is mounted correctly:

```
systemctl status glusterd.service
```

You should see that the GlusterFS volume has mounted successfully, and the service is running without issues.

This setup ensures that your GlusterFS volumes are reliably mounted after each system boot.

---

Revision #4

Created 8 September 2024 13:32:48 by aeoneros

Updated 8 September 2024 14:35:07 by aeoneros