6.2.3.2.2. Prerequisites to run NFS-Ganesha

Ensure that the following prerequisites are taken into consideration before you run NFS-Ganesha in your environment:

- A Red Hat Gluster Storage volume must be available for export and NFS-Ganesha rpms are installed.
- Only one of NFS-Ganesha, gluster-NFS or kernel-NFS servers can be enabled on a given machine/host as all NFS implementations use the port 2049 and only one can be active at a given time. Hence you must disable kernel-NFS before NFS-Ganesha is started.

  Disable the kernel-nfs using the following command:

  For Red Hat Enterprise Linux 7

  ```
  # systemctl stop nfs-server
  # systemctl disable nfs-server
  ```

  To verify if kernel-nfs is disabled, execute the following command:

  ```
  # systemctl status nfs-server
  ```

  The service should be in stopped state.

  Note

  Gluster NFS will be stopped automatically when NFS-Ganesha is enabled.

- Ensure to configure the ports as mentioned in *Port/Firewall Information for NFS-Ganesha*.
- Edit the ganesha-ha.conf file based on your environment.
- Reserve virtual IPs on the network for each of the servers configured in the ganesha.conf file. Ensure that these IPs are different than the hosts' static IPs and are not used anywhere else in the trusted storage pool or in the subnet.
- Ensure that all the nodes in the cluster are DNS resolvable. For example, you can populate the /etc/hosts with the details of all the nodes in the cluster.
- Make sure the SELinux is in **`Enforcing`** mode.
- On Red Hat Enterprise Linux 7, execute the following commands to disable and stop NetworkManager service and to enable the network service.

  ```
  # systemctl disable NetworkManager
  # systemctl stop NetworkManager
  # systemctl enable network
  ```

- Start network service on all machines using the following command:

  For Red Hat Enterprise Linux 7:

  ```
  # systemctl start network
  ```

- Create and mount a gluster shared volume by executing the following command:

```
# gluster volume set all cluster.enable-shared-storage enable
volume set: success
```

  For more information, see Section 11.8, “Setting up Shared Storage Volume”

- Create a directory named **nfs-ganesha** under **/var/run/gluster/shared_storage**
- Copy the **ganesha.conf** and **ganesha-ha.conf** files from **/etc/ganesha** to **/var/run/gluster/shared_storage/nfs-ganesha**.
- Enable the glusterfssharedstorage.service service using the following command:

```
systemctl enable glusterfssharedstorage.service
```

- Enable the nfs-ganesha service using the following command:

```
systemctl enable nfs-ganesha
```

6.2.3.2.3. Configuring the Cluster Services

The HA cluster is maintained using Pacemaker and Corosync. Pacemaker acts a resource manager and Corosync provides the communication layer of the cluster. For more information about Pacemaker/Corosync see the documentation under the _Clustering_ section of the Red Hat Enterprise Linux 7 documentation: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/

Note

It is recommended to use 3 or more nodes to configure NFS Ganesha HA cluster, in order to maintain cluster quorum.

1. Enable the pacemaker service using the following command:

   For Red Hat Enterprise Linux 7:

```
# systemctl enable pacemaker.service
```

2. Start the pcsd service using the following command.

   For Red Hat Enterprise Linux 7:

```
# systemctl start pcsd
```