

→ **‘GlusterFS’** → **Replicación de Volúmenes** → **‘HA’**.

→ **Notaciones.**

→ Indica que debe hacerse para **ambos** ‘Servidores’.

[root@srv1 ~]# → Solo para el **srv1**.

[root@srv2 ~]# → Solo para el **srv2**.

→ **Dependencia de artículo precedente:** →

<https://www.cadilinea.com/blog/wp-content/uploads/2020/03/01-RedHat-Cluster-Basico-v3.0.pdf>

→ **Objetivo inicial de montaje para los ‘Filesystem’s’.**

```
# lsblk
NAME      MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda       8:0  0  40G  0 disk
├─sda1    8:1  0   1G  0 part /boot
├─sda2    8:2  0  39G  0 part
│ └─cl-root 253:0  0  36G  0 lvm /
│ └─cl-swap 253:1  0  3,1G  0 lvm [SWAP]
sdb       8:16  0  40G  0 disk
sr0       11:0  1   7G  0 rom
```

→ **srv1**

```
...
/dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01 xfs defaults 1 2
srv1.enermol.lan:/repVol-01 /srv1/repVol-01/ glusterfs _netdev 1 2
srv2:/repVol-01 /srv2/nfsganesha nfs _netdev 1 2
...
```

→ **srv2**

```
...
/dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02 xfs defaults 1 2
srv2.enermol.lan:/repVol-01 /srv2/repVol-01/ glusterfs _netdev 1 2
srv1:/repVol-01 /srv1/nfsganesha nfs _netdev 1 2
...
```

→ **Preparación del entorno ‘HA’.**

```
# getenforce
Enforcing
```

```
# dnf install setroubleshoot setroubleshoot-server -y
```

```
# cat /etc/hosts
```

```
#127.0.0.1  srv1  srv1
```

```
#127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
#::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
```

```
192.168.10.150 pgsq1-vip.enermol.lan pgsq1-vip
192.168.10.160 pgsq1-alt-vip.enermol.lan pgsq1-alt-vip

192.168.10.151 srv1.enermol.lan srv1
192.168.10.152 srv2.enermol.lan srv2

192.168.10.161 srv1-alt.enermol.lan srv1-alt
192.168.10.162 srv2-alt.enermol.lan srv2-alt

192.168.10.159 ilo-srv1.enermol.lan ilo-srv1
192.168.10.169 ilo-srv2.enermol.lan ilo-srv2
```

→ Preparación de los ‘dispositivos de bloque’.

```
# dnf install lvm2* -y
```

```
# pvs && vgs && lvs
```

```
PV      VG Fmt Attr PSize  PFree
/dev/sda2 cl lvm2 a-- <39,00g  0
VG #PV #LV #SN Attr  VSize  VFree
cl  1  2  0 wz--n- <39,00g  0
LV   VG Attr   LSize Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
root cl -wi-ao---- 35,94g
swap cl -wi-ao----  3,05g
```

```
# pvcreate /dev/sdb
```

```
Physical volume "/dev/sdb" successfully created.
```

```
# vgcreate vg_bricks /dev/sdb
```

```
Volume group "vg_bricks" successfully created
```

```
[root@srv1 ~]# lvcreate -n lv_brickpool-01 -l 100%FREE vg_bricks
```

```
[root@srv2 ~]# lvcreate -n lv_brickpool-02 -l 100%FREE vg_bricks
```

```
[root@srv1 ~]# pvs && vgs && lvs
```

```
PV      VG      Fmt Attr PSize  PFree
/dev/sda2 cl      lvm2 a-- <39,00g  0
/dev/sdb vg_bricks lvm2 a-- <40,00g  0
VG      #PV #LV #SN Attr  VSize  VFree
cl      1  2  0 wz--n- <39,00g  0
vg_bricks 1  1  0 wz--n- <40,00g  0
LV      VG      Attr   LSize Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
root    cl      -wi-ao---- 35,94g
swap    cl      -wi-ao----  3,05g
```

```
lv_brickpool-01 vg_bricks -wi-a----- <40,00g
```

```
[root@srv2 ~]# pvs && vgs && lvs
```

```
PV      VG      Fmt Attr PSize  PFree
/dev/sda2 cl      lvm2 a-- <39,00g  0
/dev/sdb vg_bricks lvm2 a-- <40,00g  0
VG      #PV #LV #SN Attr  VSize  VFree
cl      1  2  0 wz--n- <39,00g  0
vg_bricks 1  1  0 wz--n- <40,00g  0
LV      VG      Attr   LSize  Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
root    cl      -wi-ao----- 35,94g
swap    cl      -wi-ao-----  3,05g
lv_brickpool-02 vg_bricks -wi-a----- <40,00g
```

```
[root@srv1 ~]# mkfs.xfs -i size=512 /dev/mapper/vg_bricks-lv_brickpool--01
```

```
[root@srv2 ~]# mkfs.xfs -i size=512 /dev/mapper/vg_bricks-lv_brickpool--02
```

```
[root@srv1 ~]# mkdir -p /bricks/replBrick-01
```

```
[root@srv2 ~]# mkdir -p /bricks/replBrick-02
```

```
[root@srv1 ~]# lsblk
```

```
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0  0  40G  0 disk
├─sda1                               8:1  0   1G  0 part /boot
├─sda2                               8:2  0  39G  0 part
│   ├─cl-root                        253:0  0  36G  0 lvm /
│   └─cl-swap                        253:1  0  3,1G  0 lvm [SWAP]
sdb                                  8:16  0  40G  0 disk
└─vg_bricks-lv_brickpool--01 253:2  0  40G  0 lvm
sr0                                  11:0  1   7G  0 rom
```

```
[root@srv2 ~]# lsblk
```

```
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0  0  40G  0 disk
├─sda1                               8:1  0   1G  0 part /boot
├─sda2                               8:2  0  39G  0 part
│   ├─cl-root                        253:0  0  36G  0 lvm /
│   └─cl-swap                        253:1  0  3,1G  0 lvm [SWAP]
sdb                                  8:16  0  40G  0 disk
└─vg_bricks-lv_brickpool--02 253:2  0  40G  0 lvm
sr0                                  11:0  1   7G  0 rom
```

→ Montaje 'Temporal'.

```
[root@srv1 ~]# mount /dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01
```

```
[root@srv2 ~]# mount /dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02
```

```
[root@srv1 ~]# df -hT
```

S.ficheros	Tipo	Tamaño	Usados	Disp	Usado%	Montado en
devtmpfs	devtmpfs	1,4G	0	1,4G	0%	/dev
tmpfs	tmpfs	1,4G	53M	1,4G	4%	/dev/shm
tmpfs	tmpfs	1,4G	8,7M	1,4G	1%	/run
tmpfs	tmpfs	1,4G	0	1,4G	0%	/sys/fs/cgroup
<u>/dev/mapper/cl-root</u>	<u>xfs</u>	<u>36G</u>	<u>2,8G</u>	<u>34G</u>	<u>8%</u>	<u>/</u>
<u>/dev/sda1</u>	<u>ext4</u>	<u>976M</u>	<u>190M</u>	<u>720M</u>	<u>21%</u>	<u>/boot</u>
tmpfs	tmpfs	284M	0	284M	0%	/run/user/0
/dev/mapper/vg_bricks-lv_brickpool--01	xfs	40G	318M	40G	1%	/bricks/replBrick-01

```
root@srv2 ~]# df -hT
```

S.ficheros	Tipo	Tamaño	Usados	Disp	Usado%	Montado en
devtmpfs	devtmpfs	1,4G	0	1,4G	0%	/dev
tmpfs	tmpfs	1,4G	38M	1,4G	3%	/dev/shm
tmpfs	tmpfs	1,4G	8,7M	1,4G	1%	/run
tmpfs	tmpfs	1,4G	0	1,4G	0%	/sys/fs/cgroup
<u>/dev/mapper/cl-root</u>	<u>xfs</u>	<u>36G</u>	<u>2,8G</u>	<u>34G</u>	<u>8%</u>	<u>/</u>
<u>/dev/sda1</u>	<u>ext4</u>	<u>976M</u>	<u>190M</u>	<u>720M</u>	<u>21%</u>	<u>/boot</u>
tmpfs	tmpfs	284M	0	284M	0%	/run/user/0
/dev/mapper/vg_bricks-lv_brickpool--02	xfs	40G	318M	40G	1%	/bricks/replBrick-02

```
→ Instalar 'glusterFS'.
```

```
# dnf install centos-release-gluster -y
```

```
# dnf update -y
```

```
# dnf provides glusterfs | grep -i 7
```

```
glusterfs-7.2-1.el8.x86_64 : Distributed File System  
Repositorio : centos-gluster7-test  
Proporciona : glusterfs = 7.2-1.el8  
glusterfs-7.3-1.el8.x86_64 : Distributed File System  
Repositorio : centos-gluster7-test  
Proporciona : glusterfs = 7.3-1.el8  
glusterfs-7.4-1.el8.x86_64 : Distributed File System  
Proporciona : glusterfs = 7.4-1.el8  
glusterfs-7.4-1.el8.x86_64 : Distributed File System  
Repositorio : centos-gluster7-test  
Proporciona : glusterfs = 7.4-1.el8
```

```
# ls /etc/yum.repos.d/
```

```
CentOS-AppStream.repo  CentOS-CR.repo      CentOS-Extras.repo   CentOS-HA.repo  
CentOS-Sources.repo   epel-modular.repo   epel-testing-modular.repo  
CentOS-Base.repo      CentOS-Debuginfo.repo CentOS-fasttrack.repo CentOS-Media.repo  
CentOS-Storage-common.repo epel-playground.repo epel-testing.repo  
CentOS-centosplus.repo CentOS-Devel.repo    CentOS-Gluster-7.repo CentOS-  
PowerTools.repo      CentOS-Vault.repo    epel.repo
```

```
# vim /etc/yum.repos.d/CentOS-Gluster-7.repo
# CentOS-Gluster-7.repo
#
# Please see http://wiki.centos.org/SpecialInterestGroup/Storage for more
# information
```

[centos-gluster7]

```
name=CentOS-$releasever - Gluster 7
mirrorlist=http://mirrorlist.centos.org?arch=\$basearch&release=\$releasever&repo=storage-gluster-7
#baseurl=http://mirror.centos.org/\$contentdir/\$releasever/storage/\$basearch/gluster-7/
gpgcheck=1
enabled=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-SIG-Storage
```

[centos-gluster7-test]

```
name=CentOS-$releasever - Gluster 7 Testing
baseurl=http://buildlogs.centos.org/centos/\$releasever/storage/\$basearch/gluster-7/
gpgcheck=0
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-SIG-Storage
```

dnf repolist

Última comprobación de caducidad de metadatos hecha hace 0:00:48, el mié 25 mar 2020 13:41:00 CET.

id del repositorio	nombre del repositorio
estado	
AppStream	CentOS-8 - AppStream
5.107	
BaseOS	CentOS-8 - Base
2.110	
centos-gluster7	CentOS-8 - Gluster 7
17	
*epel	Extra Packages for Enterprise Linux 8 -
x86_64	5.074
*epel-modular	Extra Packages for Enterprise Linux
Modular 8 - x86_64	0
extras	CentOS-8 - Extras

dnf update -y

dnf -y install yum-utils -y

yum-config-manager --enable PowerTools

dnf update -y

dnf repolist

```
CentOS-8 - AppStream 5.4 kB/s | 4.3 kB 00:00
```

```
CentOS-8 - Base                7.4 kB/s | 3.8 kB   00:00
CentOS-8 - Extras              2.6 kB/s | 1.5 kB   00:00
CentOS-8 - Gluster 7          3.9 kB/s | 3.0 kB   00:00
CentOS-8 - PowerTools          93 kB/s | 2.0 MB    00:22
Extra Packages for Enterprise Linux Modular 8 - x86_64      22 kB/s | 32 kB    00:01
Extra Packages for Enterprise Linux 8 - x86_64              27 kB/s | 27 kB    00:01
GlusterFS is a clustered file-system capable of scaling to several petabytes. 2.9 kB/s | 3.0 kB
00:01
GlusterFS is a clustered file-system capable of scaling to several petabytes. 3.7 kB/s | 3.0 kB
00:00
id del repositorio    nombre del repositorio                estado
AppStream             CentOS-8 - AppStream                  5.107
BaseOS                CentOS-8 - Base                       2.110
PowerTools            CentOS-8 - PowerTools                 1.523
centos-gluster7      CentOS-8 - Gluster 7                  17
epel                  Extra Packages for Enterprise Linux 8 - x86_64      5.074
epel-modular          Extra Packages for Enterprise Linux Modular 8 - x86_64  0
extras                CentOS-8 - Extras                      9
glusterfs-noarch-rhel8  GlusterFS is a clustered file-system capable of scaling to several petabytes.
1
glusterfs-rhel8      GlusterFS is a clustered file-system capable of scaling to several petabytes.
29
```

```
# dnf install glusterfs-* -y
```

```
# systemctl enable --now glusterd && systemctl status glusterd
```

```
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled; vendor preset: enabled)
   Active: active (running) since Wed 2020-03-25 16:36:43 CET; 4s ago
     Docs: man:glusterd(8)
   Process: 10279 ExecStart=/usr/sbin/glusterd -p /var/run/glusterd.pid --log-level $LOG_LEVEL
$GLUSTERD_OPTION>
   Main PID: 10280 (glusterd)
     Tasks: 9 (limit: 17950)
    Memory: 4.1M
   CGroup: /system.slice/glusterd.service
           └─10280 /usr/sbin/glusterd -p /var/run/glusterd.pid --log-level INFO
```

```
mar 25 16:36:42 srv1.enermol.lan systemd[1]: Starting GlusterFS, a clustered file-system server...
mar 25 16:36:43 srv1.enermol.lan systemd[1]: Started GlusterFS, a clustered file-system server.
```

```
# gluster --version
```

```
glusterfs 7.4
```

```
Repository revision: git://git.gluster.org/glusterfs.git
Copyright (c) 2006-2016 Red Hat, Inc. <https://www.gluster.org/>
GlusterFS comes with ABSOLUTELY NO WARRANTY.
It is licensed to you under your choice of the GNU Lesser
```

General Public License, version 3 or any later version (LGPLv3 or later), or the GNU General Public License, version 2 (GPLv2), in all cases as published by the Free Software Foundation.

→ **‘Peer’s’** → **‘GlusterFS’**.

```
# firewall-cmd --permanent --add-service=glusterfs
# firewall-cmd --reload
```

```
[root@srv1 ~]# gluster peer probe srv2.enermol.lan
peer probe: success.
```

```
[root@srv2 ~]# gluster peer probe srv1.enermol.lan
peer probe: success.
```

```
[root@srv1 ~]# lsblk -f
```

NAME	FSTYPE	LABEL	UUID	MOUNTPOINT
sda				
├─sda1	ext4		73729d53-b304-4429-8576-443befb733b1	/boot
└─sda2	LVM2_member		xcCLpn-utrQ-TTfd-rV5H-V1kG-gYCP-B0Kq00	
├─cl-root	xfs		e2a18af2-c5c9-489c-b787-35569c948620	/
└─cl-swap	swap		076f3fb1-44e8-4934-b275-0167f09b35dc	[SWAP]
sdb	LVM2_member		DAzhzk-HLxq-9OWn-ka3X-ftKH-gHYO-GVtdCQ	
├─vg_bricks-lv_brickpool--01	xfs		a37f49ac-2b03-4d45-8619-aa4bc3f64f69	/bricks/replBrick-01
sr0	iso9660	CentOS-8-1-1911-x86_64-dvd	2020-01-03-21-42-40-00	

```
[root@srv2 ~]# lsblk -f
```

NAME	FSTYPE	LABEL	UUID	MOUNTPOINT
sda				
├─sda1	ext4		1d61ef2d-6a6b-4a0a-9a3e-bfd8c60a6508	/boot
└─sda2	LVM2_member		u71JaU-5erm-8p3s-LKZd-8Ku6-Doyi-tIqS46	
├─cl-root	xfs		2a6bc616-3109-416d-908a-d1c48be57c0f	/
└─cl-swap	swap		a93cf46c-0cc6-4acf-a6cf-81d7aeb817ba	[SWAP]
sdb	LVM2_member		HS1jQi-NQZE-qA70-QI51-1b8Q-a96t-djeoR6	
├─vg_bricks-lv_brickpool--02	xfs		bd427755-99e6-44c1-b5c5-6e57afd5bcf8	/bricks/replBrick-02
sr0	iso9660	CentOS-8-1-1911-x86_64-dvd	2020-01-03-21-42-40-00	

```
[root@srv1 ~]# mkdir -p /bricks/replBrick-01/data
```

```
[root@srv2 ~]# mkdir -p /bricks/replBrick-02/data
```

```
[root@srv1 ~]# gluster volume create repVol-01 replica 2 \
srv1.enermol.lan:/bricks/replBrick-01/data \
srv2.enermol.lan:/bricks/replBrick-02/data
```

> **Nos Indica a continuación que debe existir un árbitro siempre que existan mas de ‘3’**

Volúmenes ! → No es el caso > Solo tenemos '2'!. En Producción un mínimo de 3 Servers !!.

Replica 2 volumes are prone to split-brain. **Use Arbiter or Replica 3 to avoid this.** See: <http://docs.gluster.org/en/latest/Administrator%20Guide/Split%20brain%20and%20ways%20to%20deal%20with%20it/>.

Do you still want to continue?

(y/n) y

volume create: **repVol-01**: success: please start the volume to access data

```
[root@srv1 ~]# gluster volume start repVol-01
```

```
volume start: repVol-01: success
```

```
# gluster volume info repVol-01
```

Volume Name: **repVol-01**

Type: Replicate

Volume ID: a6071a2f-c545-4713-a0e1-69f7ed1696fa

Status: Started

Snapshot Count: 0

Number of Bricks: 1 x 2 = 2

Transport-type: tcp

Bricks:

Brick1: srv1.enermol.lan:/bricks/replBrick-01/data

Brick2: srv2.enermol.lan:/bricks/replBrick-02/data

Options Reconfigured:

transport.address-family: inet

nfs.disable: on

performance.client-io-threads: off

→ **Montamos en modo: → Cliente → 'De forma Temporal'.**

```
[root@srv1 ~]# mkdir -p /srv1/repVol-01
```

```
[root@srv2 ~]# mkdir -p /srv2/repVol-01
```

```
[root@srv1 ~]# mount.glusterfs srv1.enermol.lan:/repVol-01 /srv1/repVol-01
```

```
[root@srv2 ~]# mount.glusterfs srv2.enermol.lan:/repVol-01 /srv2/repVol-01
```

```
[root@srv1 ~]# echo "En => srv1 -----" && df -hT && echo "En => srv2 -----" && ssh root@srv2 'df -hT'
```

En => srv1 -----

S.ficheros	Tipo	Tamaño Usados	Disp	Uso%	Montado en
devtmpfs	devtmpfs	1,4G 0	1,4G	0%	/dev
tmpfs	tmpfs	1,4G 38M	1,4G	3%	/dev/shm
tmpfs	tmpfs	1,4G 8,6M	1,4G	1%	/run
tmpfs	tmpfs	1,4G 0	1,4G	0%	/sys/fs/cgroup


```

/dev/mapper/cl-root          xfs      36G  2,6G  34G  8% /
/dev/sda1                    ext4     976M 190M 720M 21% /boot
tmpfs                        tmpfs    284M  0 284M  0% /run/user/0
/dev/mapper/vg_bricks-lv_brickpool--01 xfs      40G  318M  40G  1% /bricks/replBrick-01
srv1.enermol.lan:/repVol-01 fuse.glusterfs 40G  728M  40G  2% /srv1/repVol-01
En => srv2 -----
S.ficheros                   Tipo    Tamaño Usados  Disp Uso% Montado en
devtmpfs                     devtmpfs 1,4G   0 1,4G  0% /dev
tmpfs                         tmpfs    1,4G  53M 1,4G  4% /dev/shm
tmpfs                         tmpfs    1,4G  8,6M 1,4G  1% /run
tmpfs                         tmpfs    1,4G   0 1,4G  0% /sys/fs/cgroup
/dev/mapper/cl-root          xfs      36G  2,5G  34G  7% /
/dev/sda1                    ext4     976M 190M 720M 21% /boot
tmpfs                        tmpfs    284M  0 284M  0% /run/user/0
/dev/mapper/vg_bricks-lv_brickpool--02 xfs      40G  318M  40G  1% /bricks/replBrick-02
srv2.enermol.lan:/repVol-01 fuse.glusterfs 40G  728M  40G  2% /srv2/repVol-01

```

→ **‘Pruebas cruzadas’.**

```

[root@srv1 ~]# touch /srv1/repVol-01/Desde_srv1.test
[root@srv2 ~]# touch /srv2/repVol-01/Desde_srv2.test

```

```

[root@srv1 ~]# ls /srv1/repVol-01/
[root@srv2 ~]# ls /srv2/repVol-01/
Desde_srv1.test Desde_srv2.test

```

shutdown -r 0

→ **De Forma ‘Permanente’.**

```

[root@srv1 ~]# vim /etc/fstab
...
/dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01 xfs defaults 1 2
srv1.enermol.lan:/repVol-01 /srv1/repVol-01/ glusterfs _netdev 1 2
...

```

```

[root@srv2 ~]# vim /etc/fstab
...
/dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02 xfs defaults 1 2
srv2.enermol.lan:/repVol-01 /srv2/repVol-01/ glusterfs _netdev 1 2
...

```

mount -a

```

[root@srv1 ~]# df -hT
S.ficheros                   Tipo    Tamaño Usados  Disp Uso% Montado en
devtmpfs                     devtmpfs 1,4G   0 1,4G  0% /dev

```

```
tmpfs                tmpfs                1,4G  53M  1,4G  4% /dev/shm
tmpfs                tmpfs                1,4G  8,7M  1,4G  1% /run
tmpfs                tmpfs                1,4G   0  1,4G  0% /sys/fs/cgroup
/dev/mapper/cl-root  xfs                  36G   2,8G  34G  8% /
/dev/sda1            ext4                 976M  190M  720M  21% /boot
/dev/mapper/vg_bricks-lv_brickpool--01 xfs                  40G   318M  40G  1% /bricks/replBrick-01
tmpfs                tmpfs                284M   0  284M  0% /run/user/0
srv1.enermol.lan:/repVol-01 fuse.glusterfs      40G   728M  40G  2% /srv1/repVol-01
```

```
[root@srv2 ~]# df -hT
```

```
S.ficheros          Tipo      Tamaño Usados  Disp Uso% Montado en
devtmpfs            devtmpfs  1,4G   0  1,4G  0% /dev
tmpfs               tmpfs     1,4G   38M  1,4G  3% /dev/shm
tmpfs               tmpfs     1,4G   8,6M  1,4G  1% /run
tmpfs               tmpfs     1,4G   0  1,4G  0% /sys/fs/cgroup
/dev/mapper/cl-root xfs       36G   2,8G  34G  8% /
/dev/sda1           ext4      976M  190M  720M  21% /boot
/dev/mapper/vg_bricks-lv_brickpool--02 xfs       40G   318M  40G  1% /bricks/replBrick-02
tmpfs               tmpfs     284M   0  284M  0% /run/user/0
srv2.enermol.lan:/repVol-01 fuse.glusterfs 40G   728M  40G  2% /srv2/repVol-01
```

```
[root@srv1 ~]# ls /srv1/repVol-01/
```

```
Desde-srv1.txt Desde-srv2.txt
```

```
[root@srv2 ~]# ls /srv2/repVol-01/
```

```
Desde-srv1.txt Desde-srv2.txt
```

→ **splitbrain** → Consultar Documentación al respecto → **IMPORTANTE !!**.

<https://docs.gluster.org/en/latest/Troubleshooting/resolving-splitbrain/>

→ **Instalar** → 'NFS Ganesha' → `srv{1,2}`

```
# firewall-cmd --permanent --add-service=nfs
```

```
# firewall-cmd --reload
```

```
# getenforce
```

```
Enforcing
```

```
# gluster volume info
```

```
Volume Name: repVol-01
```

```
Type: Replicate
```

```
Volume ID: d9301cec-03ff-4c0c-b091-9347b688f9ed
```

```
Status: Started
```

```
Snapshot Count: 0
```

```
Number of Bricks: 1 x 2 = 2
```

```
Transport-type: tcp
```

Bricks:

Brick1: srv1.enermol.lan:/bricks/replBrick-01/data

Brick2: srv2.enermol.lan:/bricks/replBrick-02/data

Options Reconfigured:

performance.client-io-threads: off

nfs.disable: on

storage.fips-mode-rchecksum: on

transport.address-family: inet

gluster volume set repVol-01 nfs.disable on

gluster volume get repVol-01 nfs.disable

Option	Value
--------	-------

nfs.disable	on
--------------------	-----------

systemctl status nfs-server

● nfs-server.service - NFS server and services

Loaded: loaded (/usr/lib/systemd/system/nfs-server.service; **disabled**; vendor preset: **disabled**)

Active: inactive (dead)

dnf install centos-release-nfs-ganesha30 -y

dnf update -y

dnf repolist

Última comprobación de caducidad de metadatos hecha hace 0:00:15, el jue 26 mar 2020 13:58:38 CET.

id del repositorio	nombre del repositorio
estado	
AppStream	CentOS-8 - AppStream
5.120	
BaseOS	CentOS-8 - Base
2.126	
PowerTools	CentOS-8 - PowerTools
1.525	
centos-gluster7	CentOS-8 - Gluster 7
34	
centos-nfs-ganesha3	CentOS-8 - NFS Ganesha 3
11	
epel	Extra Packages for Enterprise Linux 8 -
x86_64	5.098
*epel-modular	Extra Packages for Enterprise Linux
Modular 8 - x86_64	0
extras	CentOS-8 - Extras

dnf install nfs-ganesha nfs-ganesha-gluster nfs-utils -y

dnf info nfs-ganesha

...

Paquetes instalados

Nombre : nfs-ganesha

Versión : 3.2

Lanzamiento : 2.el8

Arquitectura : x86_64

Tamaño : 2.4 M

Fuente : nfs-ganesha-3.2-2.el8.src.rpm

Repositorio : @System

Desde repo : centos-nfs-ganesha3

Resumen : NFS-Ganesha is a NFS Server running in user space

URL : <https://github.com/nfs-ganesha/nfs-ganesha/wiki>

Licencia : LGPLv3+

Descripción : nfs-ganesha : NFS-GANESHA is a NFS Server running in user space.
: It comes with various back-end modules (called FSALs) provided as
: shared objects to support different file systems and name-spaces.

```
# vim /etc/yum.repos.d/CentOS-NFS-Ganesha-3.repo
```

```
# CentOS-NFS-Ganesha-3.repo
```

```
#
```

```
# Please see http://wiki.centos.org/SpecialInterestGroup/Storage for more
```

```
# information
```

```
[centos-nfs-ganesha3]
```

```
name=CentOS-$releasever - NFS Ganesha 3
```

```
mirrorlist=http://mirrorlist.centos.org?arch=\$basearch&release=\$releasever&repo=storage-nfsganesha-3
```

```
#baseurl=https://mirror.centos.org/\$contentdir/\$releasever/storage/\$basearch/nfsganesha-3/
```

```
gpgcheck=1
```

```
enabled=1
```

```
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-SIG-Storage
```

```
[centos-nfs-ganesha3-test]
```

```
name=CentOS-$releasever - NFS Ganesha 3 Testing
```

```
baseurl=https://buildlogs.centos.org/centos/\$releasever/storage/\$basearch/nfsganesha-3/
```

```
gpgcheck=0
```

```
enabled=0
```

```
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-SIG-Storage
```

```
# mv /etc/ganesha/ganesha.conf /etc/ganesha/ganesha.conf.sample
```

```
[root@srv1 ~]# mkdir -p /srv2/nfsganesha
```

```
[root@srv2 ~]# mkdir -p /srv1/nfsganesha
```

```
→ Para → 'srv1'.
```

```
[root@srv1 ~]# vim /etc/ganesha/ganesha.conf
```

```
EXPORT{
```

```
Export_Id = 151 ; # Export ID unique to each export => srv1 → .151
Path = "/repVol-01"; # Path of the volume to be exported. Eg: "/test_volume"
```

```
FSAL {
    name = GLUSTER;
    hostname = "srv1.enermol.lan"; # IP of one of the nodes in the trusted pool
    volume = "repVol-01"; # Volume name. Eg: "test_volume"
}
```

```
Access_type = RW; # Access permissions
Squash = No_root_squash; # To enable/disable root squashing
Disable_ACL = TRUE; # To enable/disable ACL
Pseudo = "/repVol-01"; # NFSv4 pseudo path for this export. Eg: "/test_volume_pseudo"
Protocols = "3","4" ; # NFS protocols supported
Transports = "UDP","TCP" ; # Transport protocols supported
SecType = "sys"; # Security flavors supported
}
```

[root@srv1 ~]# gluster volume info

```
Volume Name: repVol-01
Type: Replicate
Volume ID: 7c6eb95b-2cfd-4083-aa26-127932b5d9dd
Status: Started
Snapshot Count: 0
Number of Bricks: 1 x 2 = 2
Transport-type: tcp
Bricks:
Brick1: srv1.enermol.lan:/bricks/replBrick-01/data
Brick2: srv2.enermol.lan:/bricks/replBrick-02/data
Options Reconfigured:
auth.allow: 192.168.10.*
diagnostics.count-fop-hits: on
diagnostics.latency-measurement: on
transport.address-family: inet
nfs.disable: on
performance.client-io-threads: off
```

[root@srv1 ~]# gluster volume set repVol-01 nfs.disable off

```
Gluster NFS is being deprecated in favor of NFS-Ganesha Enter "yes" to continue using Gluster
NFS (y/n) y
volume set: success
```

[root@srv1 ~]# gluster volume set repVol-01 nfs.rpc-auth-allow srv2.enermol.lan

```
volume set: success
# gluster volume set repVol-01 auth.allow 192.168.10.*
volume set: success
```

```
[root@srv1 ~]# gluster volume info
```

```
Volume Name: repVol-01
Type: Replicate
Volume ID: 7c6eb95b-2cfd-4083-aa26-127932b5d9dd
Status: Started
Snapshot Count: 0
Number of Bricks: 1 x 2 = 2
Transport-type: tcp
Bricks:
Brick1: srv1.enermol.lan:/bricks/replBrick-01/data
Brick2: srv2.enermol.lan:/bricks/replBrick-02/data
Options Reconfigured:
nfs.rpc-auth-allow: srv2.enermol.lan
auth.allow: 192.168.10.*
diagnostics.count-fop-hits: on
diagnostics.latency-measurement: on
transport.address-family: inet
nfs.disable: off
performance.client-io-threads: off
```

```
[root@srv1 ~]# systemctl restart glusterd.service
```

```
[root@srv1 ~]# setenforce 0
```

```
[root@srv1 ~]# getenforce
```

```
Permissive
```

```
[root@srv1 ~]# systemctl start nfs-ganesha.service
```

```
[root@srv1 ~]# showmount -e
```

```
Export list for srv1.enermol.lan:
/repVol-01 (everyone)
```

```
[root@srv2 ~]# mount -t nfs srv1:/repVol-01 /srv1/nfsganesha
```

```
[root@srv2 ~]# df -hT
```

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.9G	0	1.9G	0%	/dev
tmpfs	tmpfs	1.9G	38M	1.9G	2%	/dev/shm
tmpfs	tmpfs	1.9G	8.6M	1.9G	1%	/run
tmpfs	tmpfs	1.9G	0	1.9G	0%	/sys/fs/cgroup
/dev/sda1	xfs	40G	4.5G	36G	12%	/
/dev/mapper/vg_bricks-lv_brickpool--02	xfs	25G	1.1G	24G	1%	/bricks/replBrick-02
tmpfs	tmpfs	379M	0	379M	0%	/run/user/1000
srv1:/repVol-01	nfs4	25G	1.4G	24G	2%	/srv1/nfsganesha

```
[root@srv2 ~]# ls /srv1/nfsganesha
```

```
Desde-srv1.txt Desde-srv2.txt
```

→ Para → `srv2`.

```
[root@srv2 ~]# vim /etc/ganesha/ganesha.conf
```

```
EXPORT{
  Export_Id = 152 ; # Export ID unique to each export => srv2 → .152
  Path = "/repVol-01"; # Path of the volume to be exported. Eg: "/test_volume"

  FSAL {
    name = GLUSTER;
    hostname = "srv2.enermol.lan"; # IP of one of the nodes in the trusted pool
    volume = "repVol-01"; # Volume name. Eg: "test_volume"
  }

  Access_type = RW; # Access permissions
  Squash = No_root_squash; # To enable/disable root squashing
  Disable_ACL = TRUE; # To enable/disable ACL
  Pseudo = "/repVol-01"; # NFSv4 pseudo path for this export. Eg: "/test_volume_pseudo"
  Protocols = "3","4" ; # NFS protocols supported
  Transports = "UDP","TCP" ; # Transport protocols supported
  SecType = "sys"; # Security flavors supported
}
```

```
[root@srv2 ~]# gluster volume set repVol-01 nfs.disable off
```

```
Gluster NFS is being deprecated in favor of NFS-Ganesha Enter "yes" to continue using Gluster
NFS (y/n) y
volume set: success
```

```
[root@srv2 ~]# gluster volume set repVol-01 nfs.rpc-auth-allow srv1.enermol.lan
```

```
volume set: success
```

```
[root@srv2 ~]# gluster volume set repVol-01 auth.allow 192.168.10.*
```

```
volume set: success
```

```
[root@srv2 ~]# gluster volume info
```

```
Volume Name: repVol-01
Type: Replicate
Volume ID: 7c6eb95b-2cfd-4083-aa26-127932b5d9dd
Status: Started
Snapshot Count: 0
Number of Bricks: 1 x 2 = 2
Transport-type: tcp
Bricks:
Brick1: srv1.enermol.lan:bricks/replBrick-01/data
Brick2: srv2.enermol.lan:bricks/replBrick-02/data
Options Reconfigured:
nfs.rpc-auth-allow: srv1.enermol.lan
```

auth.allow: 192.168.10.*

diagnostics.count-fop-hits: on

diagnostics.latency-measurement: on

transport.address-family: inet

nfs.disable: off

performance.client-io-threads: off

[root@srv2 ~]# systemctl restart glusterd.service

[root@srv1 ~]# setenforce 0

[root@srv1 ~]# getenforce

Permissive

[root@srv2 ~]# systemctl start nfs-ganesha.service

[root@srv2 ~]# showmount -e

Export list for **srv2.enermol.lan:**

/repVol-01 (everyone)

[root@srv1 ~]# mount -t nfs srv2:/repVol-01 /srv2/nfsganesha

[root@srv1 ~]# df -hT

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.9G	0	1.9G	0%	/dev
tmpfs	tmpfs	1.9G	54M	1.8G	3%	/dev/shm
tmpfs	tmpfs	1.9G	8.6M	1.9G	1%	/run
tmpfs	tmpfs	1.9G	0	1.9G	0%	/sys/fs/cgroup
/dev/sda1	xfs	40G	4.4G	36G	11%	/
/dev/mapper/vg_bricks-lv_brickpool--01	xfs	25G	1.1G	24G	1%	/bricks/replBrick-01
tmpfs	tmpfs	379M	0	379M	0%	/run/user/1000
srv2:/repVol-01	nfs4	25G	1.4G	24G	2%	/srv2/nfsganesha

[root@srv2 ~]# df -hT

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.9G	0	1.9G	0%	/dev
tmpfs	tmpfs	1.9G	38M	1.9G	2%	/dev/shm
tmpfs	tmpfs	1.9G	8.6M	1.9G	1%	/run
tmpfs	tmpfs	1.9G	0	1.9G	0%	/sys/fs/cgroup
/dev/sda1	xfs	40G	4.5G	36G	12%	/
/dev/mapper/vg_bricks-lv_brickpool--02	xfs	25G	1.1G	24G	1%	/bricks/replBrick-02
tmpfs	tmpfs	379M	0	379M	0%	/run/user/1000
srv1:/repVol-01	nfs4	25G	1.4G	24G	2%	/srv1/nfsganesha

[root@srv1 ~]# ls /mnt/

Desde-srv1.txt Desde-srv2.txt

→ **Comprobaciones cruzadas.**


```
[root@srv1 ~]# touch /srv2/nfsganesha/Desde-srv1-Ganesha.txt
[root@srv2 ~]# touch /srv1/nfsganesha/Desde-srv2-Ganesha.txt
```

```
[root@srv1 ~]# ls /srv2/nfsganesha/
Desde-srv1-Ganesha.txt Desde-srv1.txt Desde-srv2-Ganesha.txt Desde-srv2.txt
[root@srv2 ~]# ls /srv1/nfsganesha/
Desde-srv1-Ganesha.txt Desde-srv1.txt Desde-srv2-Ganesha.txt Desde-srv2.txt
```

```
[root@srv1 ~]# ls /bricks/replBrick-01/data/
Desde-srv1-Ganesha.txt Desde_srv1.test Desde-srv2-Ganesha.txt Desde_srv2.test
```

```
[root@srv2 ~]# ls /bricks/replBrick-02/data/
Desde-srv1-Ganesha.txt Desde_srv1.test Desde-srv2-Ganesha.txt Desde_srv2.test
```

→ De Forma **'Permanente'**.

```
# systemctl start nfs-ganesha.service
```

```
[root@srv1 ~]# vim /etc/fstab
```

```
...
/dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01 xfs defaults 1 2
srv1.ensemol.lan:/repVol-01 /srv1/repVol-01 glusterfs _netdev 1 2
srv2:/repVol-01 /srv2/nfsganesha nfs _netdev 1 2
...
```

```
[root@srv2 ~]# vim /etc/fstab
```

```
...
/dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02 xfs defaults 1 2
srv2.ensemol.lan:/repVol-01 /srv2/repVol-01 glusterfs _netdev 1 2
srv1:/repVol-01 /srv1/nfsganesha nfs _netdev 1 2
...
```

→ **Objetivo: SELinux** → **'Enforcing'**.

```
[root@srv1 ~]# getenforce
```

```
Permissive
```

```
[root@srv2 ~]# getenforce
```

```
Permissive
```

→ Desde: **srv1, srv2**

```
# vi ganeshanfsd.te
```

```
# creamos política de acceso:
module ganeshanfsd 1.0;
```

```
require {
    type random_device_t;
    type portmap_port_t;
    type reserved_port_t;
    type ganesha_t;
    class capability dac_override;
    class tcp_socket name_connect;
    class chr_file getattr;
}

#===== ganesha_t =====
allow ganesha_t portmap_port_t:tcp_socket name_connect;
allow ganesha_t random_device_t:chr_file getattr;
allow ganesha_t reserved_port_t:tcp_socket name_connect;
allow ganesha_t self:capability dac_override;

# checkmodule -m -M -o ganeshanfsd.mod ganeshanfsd.te
# semodule_package --outfile ganeshanfsd.pp --module ganeshanfsd.mod
# semodule -i ganeshanfsd.pp

# getenforce
Enforcing

# shutdown -r 0
```

→ **Pruebas cruzadas en un entorno** : → **'Enforcing'** → (**touch's**).

→ **srv1**

```
[root@srv1 ~]# touch /bricks/rep1Brick-01/data/touch-Desde-Brick-01.txt
[root@srv1 ~]# touch /srv1/repVol-01/Desde-srv1-SELinux.txt
[root@srv1 ~]# touch /bricks/rep1Brick-01/data/touch-Desde-Brick-01-SELinux.txt
[root@srv1 ~]# touch /srv2/nfsganesha/Desde-srv1-NFSGanesha-SELinux.txt
[root@srv1 ~]# ls /srv2/nfsganesha/
Desde-1.txt Desde-2.txt Desde-srv1-Ganesha.txt Desde-srv1-NFSGanesha-SELinux.txt Desde-
srv1-SELinux.txt Desde_srv1.test Desde-srv2-Ganesha.txt Desde_srv2.test
```

→ **srv2**

```
[root@srv2 ~]# touch /bricks/rep1Brick-02/data/Desde_srv2.txt
[root@srv2 ~]# touch /srv2/repVol-01/Touch_desde-srv2.txt
[root@srv2 ~]# touch /srv1/nfsganesha/ToUCH-Desde-ganesha-srv2.txt
[root@srv2 ~]# ls /srv2/repVol-01/
Desde-1.txt Desde-srv1-Ganesha.txt Desde-srv1-SELinux.txt Desde-srv2-Ganesha.txt
Desde_srv2.txt Touch_desde-srv2.txt
Desde-2.txt Desde-srv1-NFSGanesha-SELinux.txt Desde_srv1.test Desde_srv2.test
ToUCH-Desde-ganesha-srv2.txt
```

<https://192.168.10.151:9090/>

The screenshot shows the SELinux configuration page in the CentOS Linux web console. The page title is "Normativa de SELinux". Below the title, there is a section for "Enforce policy:" with a toggle switch that is currently turned on. Underneath, there is a section for "Errores de control de acceso de SELinux" which displays a message: "No hay ninguna alerta de SELinux." The left sidebar contains a navigation menu with various system management options, and the "SELinux" option is currently selected.

The screenshot shows the storage configuration page in the CentOS Linux web console. The page title is "Almacenamiento > srv2/repVol-01". Below the title, there is a section for "srv2/repVol-01" with a progress bar showing "0,711 / 40,0 GiB". The left sidebar contains a navigation menu with various system management options, and the "Almacenamiento" option is currently selected.

<https://192.168.10.152:9090/>

The screenshot shows the SELinux configuration page in the Centos Linux GUI. The left sidebar contains a navigation menu with items like Sistema, Registros, Almacenamiento, and SELinux. The main content area is titled 'Normativa de SELinux' and shows 'Enforce policy' as 'On'. Below this, there is a section for 'Errores de control de acceso de SELinux' which displays the message 'No hay ninguna alerta de SELinux.'

The screenshot shows the storage configuration page in the Centos Linux GUI. The left sidebar is the same as in the previous screenshot. The main content area is titled 'Almacenamiento > srv1:/repVol-01'. It shows details for a storage volume: 'srv1:/repVol-01' with buttons for 'Desmontar', 'Editar', and 'Eliminar'. Below this, it lists 'Servidor: srv1:/repVol-01' and 'Punto de Montaje: /srv1/nfsGanesha'. A progress bar indicates the 'Tamaño' (Size) as 0,711 / 40,0 GiB.

RECURSOS:

https://access.redhat.com/documentation/en-us/red_hat_gluster_storage/3.4/
<https://www.itzgeek.com/how-tos/linux/centos-how-tos/install-and-configure-glusterfs-on-centos-7-rhel-7.html>
<https://www.linuxtechi.com/setup-glusterfs-storage-on-centos-7-rhel-7/>
<https://www.voztovoice.org/?q=node/2812>
<https://www.vultr.com/docs/how-to-install-cockpit-on-centos-7>
<https://cockpit-project.org/running#rhel>
<https://docs.gluster.org/en/latest/Troubleshooting/resolving-splitbrain/>
<https://aws-labs.com/linux-scale-nfsv4-nfs-ganesha-glusterfs/>
<https://docs.gluster.org/en/v3/Administrator%20Guide/NFS-Ganesha%20GlusterFS%20Integration/>
https://docs.oracle.com/cd/E52668_01/F14050/html/gluster-access.html
<https://www.golinuxcloud.com/glusterfs-distributed-volume-centos-rhel-8/>
https://www.server-world.info/en/note?os=CentOS_8&p=glusterfs&f=4
https://www.server-world.info/en/note?os=CentOS_7&p=glusterfs&f=8
https://staged-gluster-docs.readthedocs.io/en/release3.7.0beta1/Features/glusterfs_nfs-ganesha_integration/
<https://microdevsys.com/wp/glusterfs-configuration-and-setup-w-nfs-ganesha-for-an-ha-nfs-cluster/>
<https://wiki.centos.org/SpecialInterestGroup/Storage/NFS-Ganesha>
<https://www.cadilinea.com/blog/wp-content/uploads/2018/02/LABORATORIO-SELinux.pdf>

Creative Commons

Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA

© 2020 by carlos briso. Usted es libre de copiar, distribuir y comunicar públicamente la obra y hacer obras derivadas bajo las condiciones siguientes:

- Debe reconocer y citar al autor original.
 - No puede utilizar esta obra para fines comerciales (incluyendo su publicación, a través de cualquier medio, por entidades con fines de lucro).
 - Si altera o transforma esta obra o genera una obra derivada, sólo puede distribuir la obra generada bajo una licencia idéntica a ésta. Al reutilizar o distribuir la obra, tiene que dejar bien claro los términos de la licencia de esta obra.
- Alguna de estas condiciones puede no aplicarse si se obtiene el permiso del titular de los derechos de autor. Los derechos derivados de usos legítimos u otras limitaciones no se ven afectados por lo anterior. Licencia completa en castellano.

→ La información contenida en este documento y los derivados de éste se proporcionan tal cual son y los autores no asumirán responsabilidad alguna si el usuario o lector hace mal uso de éstos.