



'GlusterFS' como Replicación de Volúmenes, para HA.

→ Notaciones.

→ Indica que debe hacerse para ambos Masters.

[root@master-01 ~]# → Solo para el master-01.

[root@master-02 ~]# → Solo para el master-02.

```
# yum install epel-release.noarch -y
```

```
# vim /etc/hosts
```

```
192.168.10.201 master-01.example.com master-01
```

```
192.168.10.202 master-02.example.com master-02
```

```
# vim /etc/ssh/sshd_config
```

```
...
```

```
PermitRootLogin yes
```

```
PasswordAuthentication yes
```

```
...
```

```
# systemctl restart sshd.service
```

```
# ssh-keygen
```

```
[root@master-01 ~]# ssh-copy-id master-02
```

```
[root@master-02 ~]# ssh-copy-id master-01
```

```
# systemctl enable --now firewalld.service
```

↘ Sincronización Horaria.

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

```
# firewall-cmd --permanent --add-port=123/udp
```

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

```
# systemctl enable --now chronyd.service chrony-wait.service
```

```
# timedatectl set-timezone Europe/Madrid
```

```
# systemctl restart chronyd.service chrony-wait.service
```

```
# chronyc tracking
```

```
# chronyc sources
```

```
# timedatectl
```

```
# chronyc sourcestats
```

↘ Preparación de los dispositivos de bloque.

```
# yum install lvm2-2.02.185-2.el7.x86_64 -y
```

```
# lsblk
```

```
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
```

```
sda 8:0 0 40G 0 disk
```



```
└─sda1 8:1 0 40G 0 part /
sdb 8:16 0 25G 0 disk
```

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

```
[root@master-01 ~]# lvcreate -n lv_brickpool-01 -l 100%FREE vg_bricks
[root@master-02 ~]# lvcreate -n lv_brickpool-02 -l 100%FREE vg_bricks
```

```
root@master-01 ~]# lvs
LV          VG      Attr      LSize   Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
lv_brickpool-01 vg_bricks -wi-a----- <25.00g
```

```
[root@master-02 ~]# lvs
LV          VG      Attr      LSize   Pool Origin Data%  Meta%  Move Log Cpy%Sync Convert
lv_brickpool-02 vg_bricks -wi-a----- <25.00g
```

```
[root@master-01 ~]# mkfs.xfs -i size=512 /dev/mapper/vg_bricks-lv_brickpool--01
[root@master-02 ~]# mkfs.xfs -i size=512 /dev/mapper/vg_bricks-lv_brickpool--02
```

```
[root@master-01 ~]# mkdir -p /bricks/replBrick-01
[root@master-02 ~]# mkdir -p /bricks/replBrick-02
```

→ Montaje Temporal.

```
[root@master-01 ~]# mount /dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01
[root@master-02 ~]# mount /dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02
```

↘ Instalar 'glusterFS'.

```
# yum install centos-release-gluster -y
# yum update -y
# yum install glusterfs-server -y
# gluster --version
glusterfs 6.5

# systemctl enable --now glusterd.service
```

↘ Peer's GlusterFS.

→ Firewalld.

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

```
[root@master-01 ~]# gluster peer probe master-02
peer probe: success.
```



```
[root@master-02 ~]# gluster peer probe master-01
peer probe: success.
```

```
[root@master-01 ~]# lsblk
NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0  0 40G  0 disk
└─sda1                8:1  0 40G  0 part /
sdb                  8:16  0 25G  0 disk
└─vg_bricks-lv_brickpool--01 253:0  0 25G  0 lvm  /bricks/replBrick-01
```

```
[root@master-02 ~]# lsblk
NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0  0 40G  0 disk
└─sda1                8:1  0 40G  0 part /
sdb                  8:16  0 25G  0 disk
└─vg_bricks-lv_brickpool--02 253:0  0 25G  0 lvm  /bricks/replBrick-02
```

```
[root@master-01 ~]# mkdir -p /bricks/replBrick-01/data
[root@master-02 ~]# mkdir -p /bricks/replBrick-02/data
```

```
[root@master-01 ~]# gluster volume create repVol-01 replica 2 master-01.example.com:/bricks/
replBrick-01/data master-02.example.com:/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@master-01 ~]# 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: master-01.example.com:/bricks/replBrick-01/data

Brick2: master-02.example.com:/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@master-01 ~]# mkdir -p /mnt/repVol-01  
[root@master-02 ~]# mkdir -p /mnt/repVol-01
```

```
[root@master-01 ~]# mount.glusterfs master-01.example.com:/repVol-01 /mnt/repVol-01  
[root@master-02 ~]# mount.glusterfs master-02.example.com:/repVol-01 /mnt/repVol-01
```

→ De Forma Permanente.

```
[root@master-01 ~]# vim /etc/fstab
```

```
...  
/dev/mapper/vg_bricks-lv_brickpool--01 /bricks/replBrick-01 xfs defaults 1 2  
master-01.example.com:/repVol-01 /mnt/repVol-01/ glusterfs _netdev 1 2  
...
```

```
[root@master-02 ~]# vim /etc/fstab
```

```
...  
/dev/mapper/vg_bricks-lv_brickpool--02 /bricks/replBrick-02 xfs defaults 1 2  
master-02.example.com:/repVol-01 /mnt/repVol-01/ glusterfs _netdev 1 2  
...
```

```
# mount -a
```

```
[root@master-01 ~]# df -hT
```

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.9G	0	1.9G	0%	/dev
tmpfs	tmpfs	1.9G	0	1.9G	0%	/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	3.6G	37G	9%	/
tmpfs	tmpfs	379M	0	379M	0%	/run/user/1000
/dev/mapper/vg_bricks-lv_brickpool--01	xfs			25G	33M	25G 1% /bricks/replBrick-01
tmpfs	tmpfs	379M	0	379M	0%	/run/user/0
master-01.example.com:/repVol-01	fuse.glusterfs			25G	289M	25G 2% /mnt/repVol-01

```
[root@master-02 ~]# df -hT
```

Filesystem	Type	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	1.9G	0	1.9G	0%	/dev
tmpfs	tmpfs	1.9G	0	1.9G	0%	/dev/shm
tmpfs	tmpfs	1.9G	8.7M	1.9G	1%	/run
tmpfs	tmpfs	1.9G	0	1.9G	0%	/sys/fs/cgroup



```
/dev/sda1          xfs          40G 3.5G 37G 9% /
tmpfs             tmpfs        379M  0 379M 0% /run/user/1000
/dev/mapper/vg_brickpool--02 xfs          25G 33M 25G 1% /bricks/replBrick-02
tmpfs             tmpfs        379M  0 379M 0% /run/user/0
master-02.example.com:/repVol-01 fuse.glusterfs 25G 289M 25G 2% /mnt/repVol-01
```

```
[root@master-01 ~]# touch /mnt/repVol-01/Desde-Master01.txt
[root@master-02 ~]# touch /mnt/repVol-01/Desde-Master02.txt
[root@master-01 ~]# ls /mnt/repVol-01/
Desde-Master01.txt Desde-Master02.txt
```

```
[root@master-02 ~]# ls /mnt/repVol-01/
Desde-Master01.txt Desde-Master02.txt
```

↘ splitbrain

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

↘ Instalar cockpit

```
# yum install cockpit -y
# yum install cockpit-storaged.noarch -y
```

→ Firewalld.

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


```
# systemctl enable --now cockpit.socket
```

```
# ss -nl |grep 9090
tcp LISTEN 0 128 [::]:9090 [::]:*
```



<https://192.168.10.201:9090/>

<https://192.168.10.202:9090/>



CENTOS LINUX root

master-01.exa...

Sistema

Registros

Almacenamiento

Redes

Cuentas

Servicios

Diagnostic Reports

SELinux

Terminal

Volcado del Kernel

KB/s Leyendo

KB/s Escribiendo

Nombre	Punto de Montaje	Tamaño
/dev/sda1	/	3,56 / 40,0 GiB
/dev/vg_bricksvlv_brickpool-01	/bricks/rep1brick-01	0,0316 / 25,0 GiB

Montajes NFS

Sin ajustes de montaje NFS

Bitácoras de Almacenamiento

```

17:31 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'block-lvm2'                                udisksd
17:31 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'physical-volume'                            udisksd
17:31 g_object_notify: object class 'UDisksLinuxLogicalVolumeObject' has no property named 'logical-volume'                        udisksd
17:31 g_object_notify: object class 'UDisksLinuxVolumeGroupObject' has no property named 'volume-group'                          udisksd
17:31 g_object_notify: object class 'UDisksObjectSkeleton' has no property named 'manager-lvm2'                               udisksd
17:31 g_object_notify: object class 'UDisksObjectSkeleton' has no property named 'manager-iscsi-initiator'                       udisksd
17:31 Loading module libudisks2_lvm2.so...
17:31 Loading module libudisks2_iscsi.so...
17:31 Acquired the name org.freedesktop.UDisks2 on the system message bus
17:31 udisks daemon version 2.7.3 starting

```

Dispositivos RAID

No se ha fijado almacenamiento como RAID

Grupos de Volúmenes

vg_bricksvlv
25,0 GiB

Discos

VBOX HARDISK (V84c2da22a-e390ecae)
40 GiB Disco Duro
R: 0 B/s W: 0 B/s

VBOX HARDISK (V888d184bd-69c524aa)
25 GiB Disco Duro
R: 0 B/s W: 0 B/s

<https://192.168.10.201:9090/users>

© carlos briso

Tinkering Around Linux → <https://www.cadilinea.com/blog/>

Octubre-2019



CENTOS LINUX root

master-01.exe...

Sistema

Registros

Almacenamiento

Redes

Cuentas

Servicios

Diagnostic Reports

SELinux

Terminal

Volcado del Kernel

Almacenamiento > vg-bricks

Grupo de volúmenes vg_bricks Renombrar Eliminar

UUID PFK113-XHU2-CPka-Reux-qppk-s7s2-IT5Dmk

Capacity 25,0 GiB, 26,8 GiB, 26839351296 bytes

Logical Volumes Crear un nuevo Volumen Lógico

> 25,0 GiB xfs Sistema de archivos /dev/vg_bricks/lv_brickpool-01

Volúmenes Físicos

VBIOX HARDISK (VB88d184bd-69c524aa) +

25,0 GiB, 0 disponible -



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/>

Creative Commons

Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA

© 2019 by carlos briso. Usted es libre de copiar, distribuir y comunicar públicamente la obra y hacer obras derivadas bajo las condiciones siguientes:

a) Debe reconocer y citar al autor original.

b) No puede utilizar esta obra para fines comerciales (incluyendo su publicación, a través de cualquier medio, por entidades con fines de lucro.

c) 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.