



'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.

```
# 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_bricks-lv_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

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


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

```
# 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/>



Archivos del sistema

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... udisksd
17:31 Loading module libudisks2_iscsi.so... udisksd
17:31 Acquired the name org.freedesktop.UDisks2 on the system message bus udisksd
17:31 udisks daemon version 2.7.3 starting udisksd
```

Dispositivos RAID

No se ha fijado almacenamiento como RAID

Grupos de Volúmenes

- vg_bricksvlv
25,0 GiB

Discos

- VBOX_HARDDISK (VBox-2da22a-e390ecae)
40 GiB Disco Duro
R: 0 B/s W: 0 B/s
- VBOX_HARDDISK (VBox88d184bd-69c524aa)
25 GiB Disco Duro
R: 0 B/s W: 0 B/s

<https://192.168.10.201:9090/users>

↘ Instalar 'NFS Ganesha' en master-{01,02}

```
# firewall-cmd --permanent --add-service=nfs
# firewall-cmd --reload
# getenforce
Enforcing
```

```
# yum install nfs-ganesha nfs-ganesha-gluster nfs-utils -y
# mv /etc/ganesha/ganesha.conf /etc/ganesha/ganesha.conf.sample
```

Para → 'master-01'.

```
[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  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   5% /bricks/replBrick-01
master-01:/repVol-01     fuse.glusterfs 25G   1.4G 24G   6% /mnt/repVol-01
master-02:/repVol-01     fuse.glusterfs 25G   1.4G 24G   6% /mnt/repVol-01
tmpfs                     tmpfs     379M   0 379M   0% /run/user/1000
```

```
[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  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  4.6G 36G  12% /
/dev/mapper/vg_bricks-lv_brickpool--02 xfs         25G  1.1G 24G  5% /bricks/replBrick-02
master-01:repVol-01 fuse.glusterfs 25G  1.4G 24G  6% /mnt/repVol-01
tmpfs            tmpfs       379M  0 379M  0% /run/user/1000
tmpfs            tmpfs       379M  0 379M  0% /run/user/0
```

```
[root@master-01 ~]# vim /etc/ganesha/ganesha.conf
```

```
EXPORT{
```

```
  Export_Id = 201 ; # Export ID unique to each export => master-01 → .201
  Path = "/repVol-01"; # Path of the volume to be exported. Eg: "/test_volume"
```

```
  FSAL {
```

```
    name = GLUSTER;
    hostname = "master-01"; # 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@master-01 ~]# 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: master-01.example.com:/bricks/replBrick-01/data
```

```
Brick2: master-02.example.com:/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@master-01 ~]# 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@master-01 ~]# gluster volume set repVol-01 nfs.rpc-auth-allow master-02
```

volume set: success

```
[root@master-01 ~]# 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: master-01.example.com:/bricks/replBrick-01/data

Brick2: master-02.example.com:/bricks/replBrick-02/data

Options Reconfigured:

nfs.rpc-auth-allow: master-02

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@master-01 ~]# systemctl restart glusterd.service
```

```
[root@master-01 ~]# systemctl start nfs-ganesha.service
```

```
[root@master-01 ~]# showmount -e
```

Export list for master-01.example.com:

/repVol-01 (everyone)

```
[root@master-02 ~]# mount -t nfs master-01:/repVol-01 /mnt
```

```
[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	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	5% /bricks/replBrick-02



```
tmpfs          tmpfs  379M  0 379M  0% /run/user/1000
master-01:/repVol-01  nfs4   25G 1.4G 24G  6% /mnt
```

➔ Para → **‘master-02’**.

```
[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  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  5% /bricks/replBrick-01
master-01:/repVol-01 fuse.glusterfs 25G  1.4G 24G  6% /mnt/repVol-01
master-02:/repVol-01 fuse.glusterfs 25G  1.4G 24G  6% /mnt/repVol-01
tmpfs          tmpfs     379M  0 379M  0% /run/user/1000
```

```
[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  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  5% /bricks/replBrick-02
tmpfs          tmpfs     379M  0 379M  0% /run/user/1000
master-01:/repVol-01  nfs4   25G  1.4G 24G  6% /mnt
```

```
[root@master-02 ~]# vim /etc/ganesha/ganesha.conf
```

```
EXPORT{
  Export_Id = 202 ; # Export ID unique to each export => master-02 → .202
  Path = "/repVol-01"; # Path of the volume to be exported. Eg: "/test_volume"

  FSAL {
    name = GLUSTER;
    hostname = "master-02"; # 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@master-02 ~]# 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@master-02 ~]# gluster volume set repVol-01 nfs.rpc-auth-allow master-01
volume set: success
[root@master-02 ~]# 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: master-01.example.com:/bricks/replBrick-01/data
Brick2: master-02.example.com:/bricks/replBrick-02/data
Options Reconfigured:
nfs.rpc-auth-allow: master-01
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@master-02 ~]# systemctl restart glusterd.service
[root@master-02 ~]# systemctl start nfs-ganesha.service
```

```
[root@master-02 ~]# showmount -e
Export list for master-02.example.com:
/repVol-01 (everyone)
```

```
[root@master-02 ~]# mount -t nfs master-01:/repVol-01 /mnt
```

```
[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  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   5% /bricks/replBrick-01
tmpfs                      tmpfs     379M   0 379M   0% /run/user/1000
```



master-02:/repVol-01 nfs4 25G 1.4G 24G 6% /mnt

[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	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	5%	/bricks/replBrick-02
tmpfs	tmpfs	379M	0	379M	0%	/run/user/1000

master-01:/repVol-01 nfs4 25G 1.4G 24G 6% /mnt

Comprobaciones cruzadas.

[root@master-01 ~]# touch /mnt/Desde-master-01.txt

[root@master-02 ~]# touch /mnt/Desde-master-02.txt

[root@master-01 ~]# ls /mnt/

Desde-master-01.txt Desde-master-02.txt

[root@master-02 ~]# ls /mnt/

Desde-master-01.txt Desde-master-02.txt

<https://192.168.10.201:9090/>

Archivos del sistema

Nombre	Punto de Montaje	Tamaño
/dev/sda1	/	4,34 / 40,0 GIB
/dev/vg_bricks/lv_brickpool-01	/bricks/replBrick-01	1,09 / 25,0 GIB

Montajes NFS

Servidor	Punto de Montaje	Tamaño
master-02:/repVol-01	/mnt	1,34 / 25,0 GIB

Bitácoras de Almacenamiento

```

16:49 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'block-uuid'      udisksd
16:49 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'physical-volume'  udisksd
16:49 g_object_notify: object class 'UDisksLinuxLogicalVolumeObject' has no property named 'logical-volume' udisksd
16:49 g_object_notify: object class 'UDisksLinuxVolumeGroupObject' has no property named 'volume-group' udisksd
16:49 g_object_notify: object class 'UDisksObjectsSkeleton' has no property named 'manager-lvm2'      udisksd
16:49 g_object_notify: object class 'UDisksObjectsSkeleton' has no property named 'manager-iscsi-initiator' udisksd
16:49 Loading module libudisks2_lvm2.so...
16:49 Loading module libudisks2_iscsi.so...
16:49 Acquired the name org.freedesktop.UDisks2 on the system message bus
16:49 udisks daemon version 2.7.3 starting
  
```



<https://192.168.10.202:9090/>

CENTOS LINUX root

master-02.exa...

Sistema

Registros

Almacenamiento

Redes

Cuentas

Servicios

Diagnostic Reports

SELinux

Terminal

Volcado del Kernel

KB/s Leyendo

KB/s Escribiendo

Archivos del sistema

Nombre	Punto de Montaje	Tamaño
/dev/urandom	/	4,49 / 40,0 GiB
/dev/vg_bricksvlv_brickpool-02	/bricks/rep1brick-02	1,09 / 25,0 GiB

Montajes NFS

Servidor	Punto de Montaje	Tamaño
master-01 /rep1ol-01	/mnt	1,34 / 25,0 GiB

Bitácoras de Almacenamiento

```

16:50 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'block-lvm2'          udisksd
16:50 g_object_notify: object class 'UDisksLinuxBlockObject' has no property named 'physical-volume'      udisksd
16:50 g_object_notify: object class 'UDisksLinuxLogicalVolumeObject' has no property named 'logical-volume' udisksd
16:50 g_object_notify: object class 'UDisksLinuxVolumeGroupObject' has no property named 'volume-group'   udisksd
16:50 g_object_notify: object class 'UDisksObjectSkeleton' has no property named 'manager-lvm2'          udisksd
16:50 g_object_notify: object class 'UDisksObjectSkeleton' has no property named 'manager-iscsi-initiator' udisksd
16:50 Loading module libudisks2_lvm2.so...
16:50 Loading module libudisks2_iscsi.so...
16:50 Acquired the name org.freedesktop.UDisks2 on the system message bus
16:50 udisks daemon version 2.7.3 starting

```

Dispositivos RAID

No se ha fijado almacenamiento como RAID

Grupos de Volumen

vg_bricksvlv
25,0 GiB

Discos

VBOX_HARDDISK (V893c3e9bc-3e054d85)
40 GiB Disco Duro
R: 0 B/s W: 656 KiB/s

VBOX_HARDDISK (V834f6390c-574d30e4)
25 GiB Disco Duro
R: 0 B/s W: 0 B/s

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

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:

- 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.