

→ **TOPOLOGÍA PARA GFS2 (RedHat):**

Anfitrión KVM

192.168.1.250 hp.cadilinea.lan **hp** # **.1.250**

Clientes de Réplica para GFS2 (initiators). Pacemaker/Corosync

192.168.10.161 rhel8-gfs2-01.cadilinea.lan **rhel8-gfs2-01** # **.161**

192.168.10.162 rhel8-gfs2-02.cadilinea.lan **rhel8-gfs2-02** # **.162**

Server iSCSI (target).

192.168.10.170 rhel8-gfs2-target.cadilinea.lan **rhel8-gfs2-target** # **.170**

→ **Creamos un cluster básico: pacemaker/corosync ⇒ HA_gfs2**

```
# dnf install pcs pacemaker fence-agents-all -y
# firewall-cmd --permanent --add-service=high-availability # 2224/tcp
# firewall-cmd --reload
```

```
# passwd hacluster
# systemctl enable --now pcsd.service
```

```
rhel8-gfs2-01 ~ # pcs host auth rhel8-gfs2-01 rhel8-gfs2-02 -u hacluster
```

Password:

rhel8-gfs2-02: Authorized

rhel8-gfs2-01: Authorized

```
rhel8-gfs2-01 ~ # pcs cluster setup HA_gfs2 --start rhel8-gfs2-01 rhel8-gfs2-02
```

```
rhel8-gfs2-01 ~ # pcs cluster enable --all
```

```
rhel8-gfs2-01 ~ # pcs status
```

Cluster name: **HA_gfs2**

WARNINGS:

No stonith devices and stonith-enabled is not false

Cluster Summary:

- * Stack: corosync
- * Current DC: rhel8-gfs2-01 (version 2.1.0-8.el8-7c3f660707) - partition with quorum
- * Last updated: Mon Dec 20 17:51:24 2021
- * Last change: Mon Dec 20 17:49:44 2021 by hacluster via crmd on rhel8-gfs2-01
- * 2 nodes configured
- * 0 resource instances configured

Node List:

- * Online: [rhel8-gfs2-01 rhel8-gfs2-02]

Full List of Resources:

- * No resources



Daemon Status:

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

⇒ **Configuramos para el clúster** ⇒ [HA_gfs2](#) ⇒ [dlm.service](#) [lvmlckd.service](#)

```
rhel8-gfs2-01 ~ # pcs resource create dlm --group locking ocf:pacemaker:controld op monitor
interval=30s on-fail=fence
```

```
rhel8-gfs2-01 ~ # pcs resource clone locking interleave=true
```

```
rhel8-gfs2-01 ~ # pcs resource create lvmlckd --group locking ocf:heartbeat:lvmlckd op
monitor interval=30s on-fail=fence
```

```
rhel8-gfs2-01 ~ # pcs property set no-quorum-policy=freeze
```

(Si no hay quorum congelamos).

```
# dnf install fence-*
```

⇒ **DESDE EL ANFITRIÓN KVM .1.250** => Configuración fencing → [fence_xvm](#)

Solo desde el anfitrión KVM ⇒ **hp**

```
hp ~ # yum install fence-virtfd fence-virtfd-libvirt fence-virtfd-multicast
```

Creamos la key:

```
hp ~ # mkdir /etc/cluster
```

```
hp ~ # dd if=/dev/urandom of=/etc/cluster/fence_xvm.key bs=4k count=11
```

```
1+0 registros leídos
```

```
1+0 registros escritos
```

```
4096 bytes (4,1 kB, 4,0 KiB) copied, 0,00112964 s, 3,6 MB/s
```

Comprobamos los interfaces para libvirt:

```
hp ~ # firewall-cmd --get-active-zones
```

```
FedoraWorkstation
```

```
interfaces: virbr0 eno1
```

```
interfaces: br10
```

Generamos la configuración de 'fencing' y para todas las máquinas virtuales KVM:

```
hp ~ # fence_virtfd -c
```

```
Module search path [/usr/lib64/fence-virt/]:
```

```
Available backends:
```

```
libvirt 0.3
```

```
Available listeners:
```

```
vsock 0.1
```



multicast 1.2

Listener modules are responsible for accepting requests from fencing clients.

Listener module [multicast]:

The multicast listener module is designed for use environments where the guests and hosts may communicate over a network using multicast.

The multicast address is the address that a client will use to send fencing requests to fence_virt.

Multicast IP Address [225.0.0.12]:

Using ipv4 as family.

Multicast IP Port [1229]:

Setting a preferred interface causes fence_virt to listen only on that interface. Normally, it listens on all interfaces.

In environments where the virtual machines are using the host machine as a gateway, this **must** be set (typically to virbr0). Set to 'none' for no interface.

Interface [virbr0]: br10

The key file is the shared key information which is used to authenticate fencing requests. The contents of this file must be distributed to each physical host and virtual machine within a cluster.

Key File [/etc/cluster/fence_xvm.key]:

Backend modules are responsible for routing requests to the appropriate hypervisor or management layer.

Backend module [libvirt]:

The libvirt backend module is designed for single desktops or servers. Do not use in environments where virtual machines may be migrated between hosts.

Libvirt URI [qemu:///system]:

Configuration complete.



=== Begin Configuration ===

```
backends {
    libvirt {
        uri = "qemu:///system";
    }
}

listeners {
    multicast {
        port = "1229";
        family = "ipv4";
        interface = "br10";
        address = "225.0.0.12";
        key_file = "/etc/cluster/fence_xvm.key";
    }
}

fence_virt {
    module_path = "/usr/lib64/fence-virt/";
    backend = "libvirt";
    listener = "multicast";
}
```

=== End Configuration ===

Replace /etc/fence_virt.conf with the above [y/N]? y

hp ~ # cat /etc/fence_virt.conf

```
backends {
    libvirt {
        uri = "qemu:///system";
    }
}

listeners {
    multicast {
        port = "1229";
        family = "ipv4";
        interface = "br10";
        address = "225.0.0.12";
        key_file = "/etc/cluster/fence_xvm.key";
    }
}
```



```
}  
  
fence_virt {  
    module_path = "/usr/lib64/fence-virt";  
    backend = "libvirt";  
    listener = "multicast";  
}
```

Habilitamos servicios y puertos:

```
hp ~ # systemctl enable --now fence_virt.service  
hp ~ # firewall-cmd --permanent --add-port=1229/udp  
hp ~ # firewall-cmd --reload
```

Copiamos la key a los nodos:

```
hp ~ # ssh 192.168.10.161 mkdir /etc/cluster  
hp ~ # ssh 192.168.10.162 mkdir /etc/cluster  
hp ~ # scp /etc/cluster/fence_xvm.key root@192.168.10.161:/etc/cluster/  
hp ~ # scp /etc/cluster/fence_xvm.key root@192.168.10.162:/etc/cluster/
```

```
hp ~ # firewall-cmd --get-active-zones
```

```
FedoraWorkstation  
interfaces: virbr0 eno1  
libvirt
```

```
hp ~ # firewall-cmd --add-port=1229/udp --permanent --zone=libvirt  
hp ~ # firewall-cmd --add-port=1229/tcp --permanent --zone=libvirt  
hp ~ # firewall-cmd --reload  
hp ~ # firewall-cmd --list-all --zone=libvirt
```

```
libvirt (active)  
target: ACCEPT  
icmp-block-inversion: no  
interfaces: br10  
sources:  
services: dhcp dhcpv6 dns ssh tftp
```

```
protocols: icmp ipv6-icmp  
forward: no  
masquerade: no  
forward-ports:  
source-ports:  
icmp-blocks:
```



rich rules:

```
rule priority="32767" reject
```

⇒ **PODEMOS INSTALAR COMO ALTERNATIVA** → [fence_scsi](#)

⇒ **DESDE LOS CLIENTES AHORA:**

```
# firewall-cmd --add-port=1229/{tcp,udp} --permanent
# firewall-cmd --reload
```

```
# fence_xvm -o list | grep gfs
```

```
rhel8-gfs2-01          a8b5cd67-b511-4fe1-b2be-29d961933c4c on
rhel8-gfs2-02          ca41312b-0d6b-4f50-8f23-b1d79d83306e on
```

```
rhel8-gfs2-01 ~ # pcs property
```

Cluster Properties:

```
cluster-infrastructure: corosync
cluster-name: HA_gfs2
dc-version: 2.1.0-8.el8-7c3f660707
have-watchdog: false
no-quorum-policy: freeze
```

⇒ **Configuramos el fencing** → [fence_all](#)

```
rhel8-gfs2-01 ~ # pcs stonith create fence_all fence_xvm key_file="/etc/cluster/fence_xvm.key"
pcmk_host_map="rhel8-gfs2-01.cadilinea.lan:rhel8-gfs2-01,rhel8-gfs2-02.cadilinea.lan:rhel8-
gfs2-02" pcmk_host_list="rhel8-gfs2-01,rhel8-gfs2-02" pcmk_host_check=static-list
```

```
rhel8-gfs2-01 ~ # pcs stonith
```

```
* fence_all (stonith:fence_xvm): Started rhel8-gfs2-01
```

⇒ **Repositorios para RedHat / Instalación:**

```
# subscription-manager repos --enable rhel-8-for-x86_64-resilientstorage-rpms
# subscription-manager repos --enable rhel-8-for-x86_64-appstream-rpms
```

```
# dnf install lvm2-lockd gfs2-utils dlm
```

⇒ **Creamos los recursos: [dlm](#) / [lvmlockd](#):**

```
rhel8-gfs2-01 ~ # pcs resource create dlm --group locking ocf:pacemaker:controld op monitor
interval=30s on-fail=fence
```

```
rhel8-gfs2-01 ~ # pcs resource clone locking interleave=true
```

```
rhel8-gfs2-01 ~ # pcs resource create lvmlockd --group locking ocf:heartbeat:lvmlockd op
monitor interval=30s on-fail=fence
```

```
rhel8-gfs2-01 ~ # pcs status --full
```



Cluster name: **HA_gfs2**

Cluster Summary:

- * Stack: corosync
- * Current DC: rhel8-gfs2-01 (1) (version 2.1.0-8.el8-7c3f660707) - partition with quorum
- * Last updated: Mon Dec 20 19:33:21 2021
- * Last change: Mon Dec 20 19:29:27 2021 by root via cibadmin on rhel8-gfs2-01
- * 2 nodes configured
- * 5 resource instances configured

Node List:

- * Online: [rhel8-gfs2-01 (1) rhel8-gfs2-02 (2)]

Full List of Resources:

- * fence_all (stonith:fence_xvm): Started rhel8-gfs2-01
- * Clone Set: locking-clone [locking]:
 - * Resource Group: locking:0:
 - * dlm (ocf::pacemaker:controld): Started rhel8-gfs2-01
 - * lvmlockd (ocf::heartbeat:lvmlockd): Started rhel8-gfs2-01
 - * Resource Group: locking:1:
 - * dlm (ocf::pacemaker:controld): Started rhel8-gfs2-02
 - * lvmlockd (ocf::heartbeat:lvmlockd): Started rhel8-gfs2-02

Migration Summary:

Tickets:

PCSD Status:

rhel8-gfs2-01: Online
rhel8-gfs2-02: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

→ **Ordenamos los recursos** → **constraint** → **loc**

Utilizando el entrelazado (interleave=true) no es necesario de momento.
Unicamente hacemos mas pegajosos los recursos:

rhel8-gfs2-01 ~ # pcs resource defaults resource-stickiness=200

ps -ef | grep dlm

```
root    7007    1 0 17:47 ?        00:00:00 dlm_controld -s 0
root    7009    7007 0 17:47 ?        00:00:00 dlm_controld -s 0
root    7096    1 0 17:47 ?        00:00:00 lvmlockd -p /run/lvmlockd.pid -A 1 -g dlm
```

dlm_tool ls



dln lockspaces

```
name      lvm_global
id        0x12aabd2d
flags     0x00000000
change    member 2 joined 1 remove 0 failed 0 seq 1,1
members   1 2
```

→ **Utilizamos iSCSI como alternativa** → **block**

Clientes de Réplica para GFS2 (initiators).

```
192.168.10.161 rhel8-gfs2-01.cadilinea.lan  rhel8-gfs2-01  # .161
192.168.10.162 rhel8-gfs2-02.cadilinea.lan  rhel8-gfs2-02  # .162
```

Server iSCSI (target).

```
192.168.10.170 rhel8-gfs2-target.cadilinea.lan rhel8-gfs2-target # .170
```

rhel8-gfs2-target ~ # dnf install targetcli

rhel8-gfs2-target ~ # lsblk /dev/vd[b-z]

```
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdb 252:16  0 20G  0 disk
vdc 252:32  0 20G  0 disk
vdd 252:48  0 20G  0 disk
```

rhel8-gfs2-target ~ # targetcli

targetcli shell version 2.1.53

Copyright 2011-2013 by Datera, Inc and others.

For help on commands, type 'help'.

/> backstores/block create name=gfs2 dev=/dev/vdb

Created block storage object gfs2 using /dev/vdb.

/> iscsi/ create iqn.2022-01.lan.cadilinea:target-gfs2

Created target iqn.2022-01.lan.cadilinea:target-gfs2.

Created TPG 1.

Global pref auto_add_default_portal=true

Created default portal listening on all IPs (0.0.0.0), port 3260.

/> iscsi/iqn.2022-01.lan.cadilinea:target-gfs2/tpg1/acls create iqn.2022-01.lan.cadilinea:target-gfs2

Created Node ACL for iqn.2022-01.lan.cadilinea:target-gfs2

/> iscsi/iqn.2022-01.lan.cadilinea:target-gfs2/tpg1/ set attribute authentication=0

Parameter authentication is now '0'.

/> iscsi/iqn.2022-01.lan.cadilinea:target-gfs2/tpg1/ set attribute generate_node_acls=1

Parameter generate_node_acls is now '1'.

/> iscsi/iqn.2022-01.lan.cadilinea:target-gfs2/tpg1/luns create /backstores/block/gfs2

Created LUN 0.

Created LUN 0->0 mapping in node ACL iqn.2022-01.lan.cadilinea:target-gfs2

/> ls

```
o- / ..... [...]
  o- backstores ..... [...]
```



```

| o- block ..... [Storage Objects: 1]
|| o- gfs2 ..... [/dev/vdb (20.0GiB) write-thru activated]
|| o- alua ..... [ALUA Groups: 1]
|| o- default_tg_pt_gp ..... [ALUA state: Active/optimized]
| o- fileio ..... [Storage Objects: 0]
| o- pscsi ..... [Storage Objects: 0]
| o- ramdisk ..... [Storage Objects: 0]
o- iscsi ..... [Targets: 1]
| o- iqn.2022-01.lan.cadilinea:target-gfs2 ..... [TPGs: 1]
| o- tpg1 ..... [gen-acls, no-auth]
| o- acls ..... [ACLs: 1]
| | o- iqn.2022-01.lan.cadilinea:target-gfs2 ..... [Mapped LUNs: 1]
| | o- mapped_lun0 ..... [lun0 block/gfs2 (rw)]
| o- luns ..... [LUNs: 1]
| | o- lun0 ..... [block/gfs2 (/dev/vdb) (default_tg_pt_gp)]
| o- portals ..... [Portals: 1]
| o- 0.0.0.0:3260 ..... [OK]
o- loopback ..... [Targets: 0]
/> exit
Global pref auto_save_on_exit=true
Configuration saved to /etc/target/saveconfig.json

```

```

rhel8-gfs2-target ~ # systemctl enable --now target.service
rhel8-gfs2-target ~ # firewall-cmd --permanent --add-port=3260/tcp
rhel8-gfs2-target ~ # firewall-cmd --reload

```

→ Desde los clientes ahora:

```

# iscsiadm -m discovery -t st -p 192.168.10.170
192.168.10.170:3260,1 iqn.2022-01.lan.cadilinea:target-gfs2
# iscsiadm -m node -T iqn.2022-01.lan.cadilinea:target-gfs2 -p 192.168.10.170 -l
Logging in to [iface: default, target: iqn.2022-01.lan.cadilinea:target-gfs2, portal:
192.168.10.170,3260]
Login to [iface: default, target: iqn.2022-01.lan.cadilinea:target-gfs2, portal: 192.168.10.170,3260]
successful
# lsblk /dev/sda
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 20G 0 disk
# vim /etc/iscsi/initiatorname.iscsi
#InitiatorName=iqn.1994-05.com.redhat:195fe743df75
InitiatorName=iqn.2022-01.lan.cadilinea:target-gfs2

```

→ Procedemos con la creación del block compartido descubierto → /dev/sda

```

rhel8-gfs2-01 ~ # vgcreate --shared --locktype dlm shared_gfs2 /dev/sda
Physical volume "/dev/sda" successfully created.
Volume group "shared_gfs2" successfully created
VG shared_gfs2 starting dlm lockspace
Starting locking. Waiting until locks are ready...

```

**# vgs**

```
Reading VG shared_vg1 without a lock.
VG      #PV #LV #SN Attr  VSize  VFree
rhel    1  3  0 wz--n- <24,00g  4,80g
shared_gfs2  1  0  0 wz--ns <20,00g <20,00g
```

→ **En un nodo alternativo (NO PRINCIPAL) => rhel8-gfs2-02**

rhel8-gfs2-02 ~ # vgchange --lock-start shared_gfs2

```
VG shared_gfs2 starting dlm lockspace
Starting locking. Waiting until locks are ready...
```

→ **En un nodo indistinto:**

rhel8-gfs2-01 ~ # lvcreate --activate sy -L5G -n lv_shared_gfs2 shared_gfs2

```
Logical volume "lv_shared_gfs2" created.
```

rhel8-gfs2-01 ~ # mkfs.gfs2 -j2 -p lock_dlm -t HA_gfs2:gfs2-compartido /dev/mapper/shared_gfs2-lv_shared_gfs2

```
/dev/mapper/shared_gfs2-lv_shared_gfs2 is a symbolic link to /dev/dm-6
```

```
This will destroy any data on /dev/dm-6
```

```
Are you sure you want to proceed? [y/n] y
```

```
Discarding device contents (may take a while on large devices): Done
```

```
Adding journals: Done
```

```
Building resource groups: Done
```

```
Creating quota file: Done
```

```
Writing superblock and syncing: Done
```

```
Device: /dev/mapper/shared_gfs2-lv_shared_gfs2
```

```
Block size: 4096
```

```
Device size: 5,00 GB (1310720 blocks)
```

```
Filesystem size: 5,00 GB (1310718 blocks)
```

```
Journals: 2
```

```
Journal size: 32MB
```

```
Resource groups: 22
```

```
Locking protocol: "lock_dlm"
```

```
Lock table: "HA_gfs2:gfs2-compartido"
```

```
UUID: 9e3f48cc-80ca-459d-b10d-5fbf3cb46f74
```

→ **Creamos un nuevo recurso → shared_lv-gfs2 para que pacemaker lo active automáticamente (clone) , y en todos los nodos:**

```
rhel8-gfs2-01 ~ # pcs resource create shared_lv-gfs2 --group shared_gfs2 ocf:heartbeat:LVM-
activate lvname=lv_shared_gfs2 vgname=shared_gfs2 activation_mode=shared
vg_access_mode=lvmlckd
```

```
rhel8-gfs2-01 ~ # pcs resource clone shared_gfs2 interleave=true
```

→ **Configuramos las restricciones → loc → para que los recursos se inicien en orden !! :**

**=> location**

No procede en principio una configuración para **location** en una configuración entrelazada. (interleave=true).

=> order

```
rhel8-gfs2-01 ~ # pcs constraint order start locking-clone then shared_gfs2-clone
```

Adding locking-clone shared_gfs2-clone (kind: Mandatory) (Options: first-action=start then-action=start)

=> colocation

```
rhel8-gfs2-01 ~ # pcs constraint colocation add shared_gfs2-clone with locking-clone
```

=> Resúmen:

```
rhel8-gfs2-01 ~ # pcs constraint
```

Location Constraints:

Ordering Constraints:

```
start locking-clone then start shared_gfs2-clone (kind:Mandatory)
```

Colocation Constraints:

```
shared_gfs2-clone with locking-clone (score:INFINITY)
```

Ticket Constraints:

→ Preparamos el montaje automático del FS → GFS2 de ejemplo:

```
rhel8-gfs2-01 ~ # pcs resource create FS-GFS2_mount --group shared_gfs2
```

```
ocf:heartbeat:Filesystem device="/dev/shared_gfs2/lv_shared_gfs2" directory="/mnt/gfs2" fstype="gfs2" options=noatime op monitor interval=10s on-fail=fence
```

```
rhel8-gfs2-01 ~ # pcs status --full
```

Cluster name: **HA_gfs2**

Cluster Summary:

- * Stack: corosync
- * Current DC: rhel8-gfs2-01 (1) (version 2.1.0-8.el8-7c3f660707) - partition with quorum
- * Last updated: Fri Dec 31 18:48:55 2021
- * Last change: Fri Dec 31 18:44:57 2021 by root via cibadmin on rhel8-gfs2-01
- * 2 nodes configured
- * 9 resource instances configured

Node List:

- * Online: [rhel8-gfs2-01 (1) rhel8-gfs2-02 (2)]

Full List of Resources:

- * Clone Set: locking-clone [locking]:

```

* Resource Group: locking:0:
  * dlm      (ocf::pacemaker:controld):  Started rhel8-gfs2-02
  * lvmlockd (ocf::heartbeat:lvmlockd):  Started rhel8-gfs2-02
* Resource Group: locking:1:
  * dlm      (ocf::pacemaker:controld):  Started rhel8-gfs2-01
  * lvmlockd (ocf::heartbeat:lvmlockd):  Started rhel8-gfs2-01
* fence_all (stonith:fence_xvm): Started rhel8-gfs2-01
* Clone Set: shared_gfs2-clone [shared_gfs2]:
  * Resource Group: shared_gfs2:0:
    * shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-02
    * FS-GFS2_mount (ocf::heartbeat:Filesystem):  Started rhel8-gfs2-02
  * Resource Group: shared_gfs2:1:
    * shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-01
    * FS-GFS2_mount (ocf::heartbeat:Filesystem):  Started rhel8-gfs2-01

```

Migration Summary:

Tickets:

PCSD Status:

```

rhel8-gfs2-01: Online
rhel8-gfs2-02: Online

```

Daemon Status:

```

corosync: active/enabled
pacemaker: active/enabled

```

dlm_tool ls

dlm lockspaces

```

name      gfs2-compartido
id        0xa0c1379f
flags     0x00000000
change    member 2 joined 1 remove 0 failed 0 seq 1,1
members   1 2

```

```

name      lvm_global
id        0x12aabd2d
flags     0x00000000
change    member 2 joined 1 remove 0 failed 0 seq 1,1
members   1 2

```

→ **Es todo OK !!!**

dlm_tool status

```

luster nodeid 2 quorate 1 ring seq 213 213
daemon now 489 fence_pid 0
node 1 M add 21 rem 0 fail 0 fence 0 at 0 0

```



```
node 2 M add 21 rem 0 fail 0 fence 0 at 0 0
```

```
→ Suspende el FS → FS-GFS2_mount
```

```
rhel8-gfs2-01 ~ # pcs resource config FS-GFS2_mount
```

```
Resource: FS-GFS2_mount (class=ocf provider=heartbeat type=Filesystem)
```

```
Attributes: device=/dev/shared_gfs2/lv_shared_gfs2 directory=/mnt/gfs2 fstype=gfs2  
options=noatime
```

```
Operations: monitor interval=10s on-fail=fence (FS-GFS2_mount-monitor-interval-10s)  
start interval=0s timeout=60s (FS-GFS2_mount-start-interval-0s)  
stop interval=0s timeout=60s (FS-GFS2_mount-stop-interval-0s)
```

```
rhel8-gfs2-01 ~ # pcs resource status FS-GFS2_mount
```

```
* Clone Set: shared_gfs2-clone [shared_gfs2]:
```

```
* Started: [ rhel8-gfs2-01 rhel8-gfs2-02 ]
```

```
rhel8-gfs2-01 ~ # pcs resource status shared_gfs2:0
```

```
* Clone Set: shared_gfs2-clone [shared_gfs2]:
```

```
* Started: [ rhel8-gfs2-02 ]
```

```
rhel8-gfs2-01 ~ # pcs resource status shared_gfs2:1
```

```
* Clone Set: shared_gfs2-clone [shared_gfs2]:
```

```
* Started: [ rhel8-gfs2-01 ]
```

```
rhel8-gfs2-01 ~ # pcs resource disable FS-GFS2_mount
```

```
rhel8-gfs2-01 ~ # pcs resource status FS-GFS2_mount
```

```
* Clone Set: shared_gfs2-clone [shared_gfs2]:
```

```
* Resource Group: shared_gfs2:0:
```

```
* shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-02
```

```
* FS-GFS2_mount (ocf::heartbeat:Filesystem): Stopped (disabled)
```

```
* Resource Group: shared_gfs2:1:
```

```
* shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-01
```

```
* FS-GFS2_mount (ocf::heartbeat:Filesystem): Stopped (disabled)
```

```
!!! Para chequear el Sistema debe estar previamente DESMONTADO !!!
```

```
rhel8-gfs2-01 ~ # fsck.gfs2 -y /dev/mapper/shared_gfs2-lv_shared_gfs2 > /tmp/fsck-gfs2.out
```

```
rhel8-gfs2-01 ~ # cat /tmp/fsck-gfs2.out
```

```
Initializing fsck
```

```
Validating resource group index.
```

```
Level 1 resource group check: Checking if all rgrp and rindex values are good.
```

```
(level 1 passed)
```

```
Starting pass1
```

```
Reconciling bitmaps.
```

```
reconcile_bitmaps completed in 0.016s
```

```
pass1 completed in 0.039s
```

```
Starting pass1b
```

```
pass1b completed in 0.000s
```

```
Starting pass2
```

```

pass2 completed in 0.001s
Starting pass3
pass3 completed in 0.000s
Starting pass4
pass4 completed in 0.003s
Starting check_statfs
check_statfs completed in 0.000s
fsck.gfs2 complete

```

rhel8-gfs2-01 ~ # pcs resource enable --wait=100 FS-GFS2_mount

Waiting for the cluster to apply configuration changes (timeout: 100 seconds)...
resource 'FS-GFS2_mount' is running on nodes 'rhel8-gfs2-01', 'rhel8-gfs2-02'

→ **Crecimiento del FS → +3G**

rhel8-gfs2-01 ~ # lsblk /dev/sda

```

NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 20G  0 disk
└─shared_gfs2-lv_shared_gfs2 253:6   0 5G  0 lvm  /mnt/gfs2

```

rhel8-gfs2-01 ~ # lvextend --lockopt skiplv -L +3G /dev/shared_gfs2/lv_shared_gfs2

WARNING: skipping LV lock in lvmlockd.

WARNING: shared LV may require refresh on other hosts where it is active.

Size of logical volume shared_gfs2/lv_shared_gfs2 changed from 5,00 GiB (1280 extents) to 8,00 GiB (2048 extents).

Logical volume shared_gfs2/lv_shared_gfs2 successfully resized.

rhel8-gfs2-01 ~ # lsblk /dev/sda

```

NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 20G  0 disk
└─shared_gfs2-lv_shared_gfs2 253:6   0 8G  0 lvm  /mnt/gfs2

```

rhel8-gfs2-01 ~ # ssh rhel8-gfs2-02 lsblk /dev/sda

```

NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 20G  0 disk
└─shared_gfs2-lv_shared_gfs2 253:6   0 5G  0 lvm  /mnt/gfs2

```

rhel8-gfs2-01 ~ # ssh rhel8-gfs2-02 lvchange --refresh /dev/shared_gfs2/lv_shared_gfs2

rhel8-gfs2-01 ~ # lsblk /dev/sda

```

NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 20G  0 disk
└─shared_gfs2-lv_shared_gfs2 253:6   0 8G  0 lvm  /mnt/gfs2

```

rhel8-gfs2-01 ~ # ssh rhel8-gfs2-02 lsblk /dev/sda

```

NAME                MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda                  8:0    0 20G  0 disk
└─shared_gfs2-lv_shared_gfs2 253:6   0 8G  0 lvm  /mnt/gfs2

```

→ **Añadir journals -j2:**

```
rhel8-gfs2-01 ~ # gfs2_edit /dev/shared_gfs2/lv_shared_gfs2
```

```
gfs2_edit - Global File System Editor (use with extreme caution)
```

```
Block #16 (0x10) of 2097152 (0x200000) (superblock)
```

```
(p.1 of 16--Resrv)
```

```
00010000 01161970 00000001 00000000 00000000 [...p.....] sb_header.mh_magic
00010010 00000064 00000000 00000709 0000076c [...d.....]
00010020 00000000 00001000 0000000c 00000000 [.....]
00010030 00000000 00000003 00000000 00002024 [..... $]
00010040 00000000 00000000 00000000 00000000 [.....]
00010050 00000000 00000010 00000000 00004248 [.....BH]
00010060 6c6f636b 5f646c6d 00000000 00000000 [lock_dlm.....]
00010070 00000000 00000000 00000000 00000000 [.....]
00010080 00000000 00000000 00000000 00000000 [.....]
00010090 00000000 00000000 00000000 00000000 [.....]
000100a0 48415f67 6673323a 67667332 2d636f6d [HA_gfs2:gfs2-com]
000100b0 70617274 69646f00 00000000 00000000 [partido.....]
000100c0 00000000 00000000 00000000 00000000 [.....]
000100d0 00000000 00000000 00000000 00000000 [.....]
000100e0 00000000 00000000 00000000 00000000 [.....]
000100f0 00000000 00000000 00000000 00000000 [.....]
```

Copyright (C) Red Hat, Inc. 2004-2010 All rights reserved. - Mode: Hex view p

```
rhel8-gfs2-01 ~ # gfs2_edit -p jindex /dev/shared_gfs2/lv_shared_gfs2 | grep journal
```

```
3/3 [fc7745eb] 1/18 (0x1/0x12) +0: File journal0
```

```
4/4 [8b70757d] 2/8231 (0x2/0x2027) +0: File journal1
```

```
rhel8-gfs2-01 ~ # gfs2_jadd -j 2 /dev/shared_gfs2/lv_shared_gfs2
```

```
Filesystem: /dev/shared_gfs2/lv_shared_gfs2
```

```
Old journals: 2
```

```
New journals: 4
```

```
rhel8-gfs2-01 ~ # gfs2_edit -p jindex /dev/shared_gfs2/lv_shared_gfs2 | grep journal
```

```
3/3 [fc7745eb] 1/18 (0x1/0x12) +0: File journal0
```

```
4/4 [8b70757d] 2/8231 (0x2/0x2027) +0: File journal1
```

```
5/5 [127924c7] 10/17237 (0xa/0x4355) +1: File journal2
```

```
6/6 [657e1451] 14/50334 (0xe/0xc49e) +1: File journal3
```

→ Si no estuviésemos en un cluster pacemaker Suspensión/Reanudación:

```
rhel8-gfs2-01 ~ # dmsetup suspend /dev/shared_gfs2/lv_shared_gfs2
```

```
rhel8-gfs2-01 ~ # dmsetup resume /dev/shared_gfs2/lv_shared_gfs2
```

→ Resúmen:

```
rhel8-gfs2-01 ~ # pcs status --full
```

```
Cluster name: HA_gfs2
```



Cluster Summary:

- * Stack: corosync
- * Current DC: rhel8-gfs2-02 (2) (version 2.1.0-8.el8-7c3f660707) - partition with quorum
- * Last updated: Sun Jan 2 19:02:21 2022
- * Last change: Sun Jan 2 18:51:48 2022 by hacluster via crmd on rhel8-gfs2-02
- * 2 nodes configured
- * 9 resource instances configured

Node List:

- * Online: [rhel8-gfs2-01 (1) rhel8-gfs2-02 (2)]

Full List of Resources:

- * Clone Set: locking-clone [locking]:
 - * Resource Group: locking:0:
 - * dlm (ocf::pacemaker:controld): Started rhel8-gfs2-02
 - * lvmlockd (ocf::heartbeat:lvmlockd): Started rhel8-gfs2-02
 - * Resource Group: locking:1:
 - * dlm (ocf::pacemaker:controld): Started rhel8-gfs2-01
 - * lvmlockd (ocf::heartbeat:lvmlockd): Started rhel8-gfs2-01
- * fence_all (stonith:fence_xvm): Started rhel8-gfs2-01
- * Clone Set: shared_gfs2-clone [shared_gfs2]:
 - * Resource Group: shared_gfs2:0:
 - * shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-02
 - * FS-GFS2_mount (ocf::heartbeat:Filesystem): Started rhel8-gfs2-02
 - * Resource Group: shared_gfs2:1:
 - * shared_lv-gfs2 (ocf::heartbeat:LVM-activate): Started rhel8-gfs2-01
 - * FS-GFS2_mount (ocf::heartbeat:Filesystem): Started rhel8-gfs2-01

Migration Summary:

Tickets:

PCSD Status:

- rhel8-gfs2-01: Online
- rhel8-gfs2-02: Online

Daemon Status:

- corosync: active/enabled
- pacemaker: active/enabled
- pcsd: active/enabled

Creative Commons

Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA

© 2022 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.