

→ **Convenciones:**

```
# En todos los nodos como 'sudo su'.  
[root@srv1 ~]# Solo en servidor 'srv1' → como 'sudo su'.  
[root@srv2 ~]# Solo en servidor 'srv2' → como 'sudo su'.
```

### 362.1 DRBD (weight: 6)

<b>Weight</b>	6
<b>Description</b>	Candidates are expected to have the experience and knowledge to install, configure, maintain and troubleshoot DRBD devices. This includes integration with Pacemaker. DRBD configuration of version 9.0.x is covered.

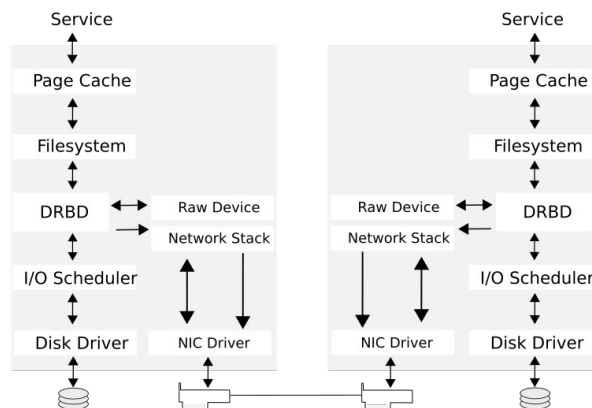
#### Key Knowledge Areas:

- Understand the DRBD architecture
- Understand DRBD resources, states and replication modes
- Configure DRBD disks and devices
- Configure DRBD networking connections and meshes
- Configure DRBD automatic recovery and error handling
- Configure DRBD quorum and handlers for split brain and fencing
- Manage DRBD using drbdadm
- Understand the principles of drbdsetup and drbdmeta
- Restore and verify the integrity of a DRBD device after an outage
- Integrate DRBD with Pacemaker
- Understand the architecture and features of LINSTOR

#### Partial list of the used files, terms and utilities:

- Protocol A, B and C
- Primary, Secondary
- Three-way replication
- drbd kernel module
- drbdadm
- drbdmon
- drbdsetup
- drbdmeta
- /etc/drbd.conf
- /etc/drbd.d/
- /proc/drbd

→ **Conceptos Teóricos previos.**



→ **Protocolos A, B y C**

**Protocol A**

Asynchronous replication protocol. Local write operations on the primary node are considered completed as soon as the local disk write has finished, and the replication packet has been placed in the local TCP send buffer. In the event of forced fail-over, data loss may occur. The data on the standby node is consistent after fail-over; however, the most recent updates performed prior to the crash could be lost. Protocol A is most often used in long distance replication scenarios. When used in combination with DRBD Proxy it makes an effective disaster recovery solution. See Long-distance replication via DRBD Proxy, for more information.

**Protocol B**

Memory synchronous (semi-synchronous) replication protocol. Local write operations on the primary node are considered completed as soon as the local disk write has occurred, and the replication packet has reached the peer node. Normally, no writes are lost in case of forced fail-over. However, in the event of simultaneous power failure on both nodes and concurrent, irreversible destruction of the primary's data store, the most recent writes completed on the primary may be lost.

**Protocol C**

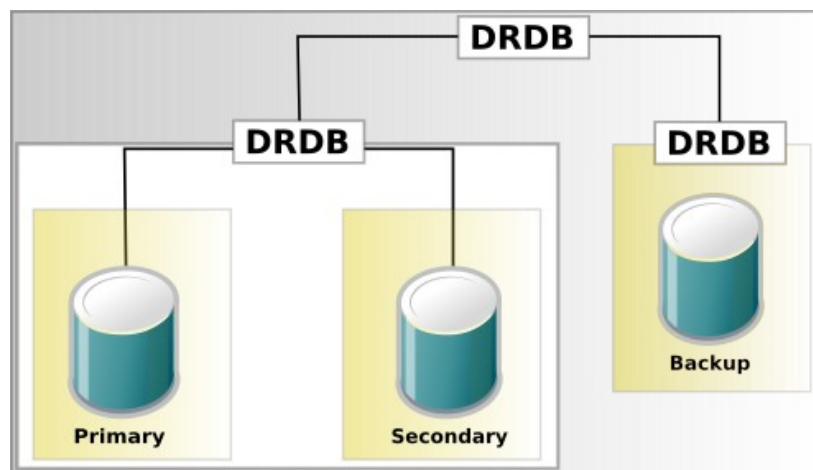
Synchronous replication protocol. Local write operations on the primary node are considered completed only after both the local and the remote disk write(s) have been confirmed. As a result,

loss of a single node is guaranteed not to lead to any data loss. Data loss is, of course, inevitable even with this replication protocol if all nodes (resp. their storage subsystems) are irreversibly destroyed at the same time.

By far, the most commonly used replication protocol in DRBD setups is protocol C.

The choice of replication protocol influences two factors of your deployment: protection and latency. Throughput, by contrast, is largely independent of the replication protocol selected.

### → Three Way Replication



**Primario → Secundario** ⇒ **Protocolo C**

**Primario → Backup** ⇒ **Protocolo A**

When using three-way replication, DRBD adds a third node to an existing 2-node cluster and replicates data to that node, where it can be used for backup and disaster recovery purposes. This type of configuration generally involves Long-distance replication via DRBD Proxy.

Three-way replication works by adding another, stacked DRBD resource on top of the existing resource holding your production data.

### → Topologia general.

# vim /etc/hosts

```
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
```

```
:::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
```

```
## hosts DRBD.
```

```
192.168.10.161 drbd-01.cadilinea.lan drbd-01
```

```
192.168.10.162 drbd-02.cadilinea.lan drbd-02
```

```
# 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
```

→ **Sincronización Horaria.**

→ → Utilizamos → **chronyd.service** → **NTP**

```
# systemctl status chronyd.service
```

- chronyd.service - NTP client/server

Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled; vendor preset: enabled)

Active: active (running) since sáb 2021-05-08 18:52:09 CEST; 7s ago

Docs: man:chronyd(8)

man:chrony.conf(5)

Process: 20994 ExecStartPost=/usr/libexec/chrony-helper update-daemon (code=exited, status=0/SUCCESS)

Process: 20990 ExecStart=/usr/sbin/chronyd \$OPTIONS (code=exited, status=0/SUCCESS)

Main PID: 20992 (chronyd)

CGroup: /system.slice/chronyd.service

└─20992 /usr/sbin/chronyd

```
may 08 18:52:09 drbd-01.cadilinea.lan systemd[1]: Stopped NTP client/server.
```

```
may 08 18:52:09 drbd-01.cadilinea.lan systemd[1]: Starting NTP client/server...
```

```
may 08 18:52:09 drbd-01.cadilinea.lan chronyd[20992]: chronyd version 3.4 starting (+CMDMON +NTP +REFCLOCK +RTC +PRIVDROP +SCFI...DEBUG)
```

```
may 08 18:52:09 drbd-01.cadilinea.lan chronyd[20992]: Frequency -13.974 +/- 0.114 ppm read
```

from /var/lib/chrony/drift

may 08 18:52:09 drbd-01.cadilinea.lan systemd[1]: Started NTP client/server.

may 08 18:52:14 drbd-01.cadilinea.lan chronyd[20992]: Selected source 178.255.228.77

may 08 18:52:16 drbd-01.cadilinea.lan chronyd[20992]: Selected source 162.159.200.1

Hint: Some lines were ellipsized, use -l to show in full.

→ [Paqueteria previa.](#)

```
# yum install lvm2-lockd gcc gcc-c++ make automake autoconf help2man libxslt libxslt-devel  
flex rpm-build kernel-devel git
```

→ [DRBD \(Linbit\).](#)

```
# rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
```

```
# yum install https://www.elrepo.org/elrepo-release-7.0-3.el7.elrepo.noarch.rpm
```

```
# yum update
```

```
# yum install drbd90-utils kmod-drbd90.x86_64 policycoreutils-python
```

```
# firewall-cmd --permanent --add-port=6996-7800/tcp
```

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

```
# modprobe drbd
```

```
# lsmod |grep drbd
```

```
drbd                573100 0
```

```
libcrc32c           12644 5 xfs,drbd,dm_persistent_data,nf_nat,nf_conntrack
```

```
# echo drbd > /etc/modules-load.d/drbd.conf
```

```
# semanage permissive -a drbd_t
```

→ [Creación de los Blocks Devices](#)

Crear partición de replica previamente → `'fdisk /dev/vdb' → /dev/vdb1'`

```
# fdisk /dev/vdb
```

Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.

Be careful before using the write command.

Device does not contain a recognized partition table

Building a new DOS disklabel with disk identifier 0x3f2c5ef5.

Orden (m para obtener ayuda): n

Partition type:

p primary (0 primary, 0 extended, 4 free)

e extended

Select (default p): **p**

Número de partición (1-4, default 1): **1**

Primer sector (2048-41943039, valor predeterminado 2048):

Se está utilizando el valor predeterminado 2048

Last sector, +sectors or +size{K,M,G} (2048-41943039, valor predeterminado 41943039):

Se está utilizando el valor predeterminado 41943039

Partition 1 of type Linux and of size 20 GiB is set

Orden (m para obtener ayuda): **w**

¡Se ha modificado la tabla de particiones!

Llamando a ioctl() para volver a leer la tabla de particiones.

Se están sincronizando los discos.

```
# pvcreate /dev/vdb1
```

```
# vgcreate vg_drbd /dev/vdb1
```

```
# lvcreate -n lv_drbd -L 1G vg_drbd
```

```
# vim /etc/drbd.d/drbd0.res
```

```
resource drbd0 {
```

```
protocol C;
```

```
meta-disk internal;
device /dev/drbd0;
disk /dev/vg_drbd/lv_drbd;
handlers {
    split-brain "/usr/lib/drbd/notify-split-brain.sh root";
}
startup {
    wfc-timeout 20;
    become-primary-on both;
}
net {
    allow-two-primaries yes;
    after-sb-0pri discard-zero-changes;
    after-sb-1pri discard-secondary;
    after-sb-2pri disconnect;
    rr-conflict disconnect;
    csums-alg sha1;
}
disk {
    on-io-error detach;
    resync-rate 10M; # 100Mbps dedicated link
    # All cluster file systems require fencing
    fencing resource-and-stonith;
}
syncer {
    verify-alg sha1;
}
on drbd-01 {
```

```
address 192.168.10.161:7789;
}
on drbd-02 {
address 192.168.10.162:7789;
}
}
```

→ Creamos los meta-datos para el recurso denominado → **drbd0**

Los recursos pueden crearse y administrarse de una forma mas sencilla con → **drbdmanage**

**# drbdadm create-md drbd0**

```
--== Thank you for participating in the global usage survey ==--
```

The server's response is:

you are the 39880th user to install this version

initializing activity log

initializing bitmap (32 KB) to all zero

Writing meta data...

New drbd meta data block successfully created.

success

**# modprobe drbd**

**# drbdadm up drbd0**

**# drbdadm dstate drbd0**

Inconsistent/DUnknown

**drbd-01 ~ # drbdadm primary drbd0 --force**

**drbd-01 ~ # drbdadm dstate drbd0**

UpToDate/UpToDate

**drbd-02 ~ # drbdadm primary drbd0 --force**

**drbd-01 ~ # drbdadm status**

drbd0 role:Primary

disk:UpToDate



drbd-02 role:Primary

peer-disk:UpToDate

# drbdmon → Multiplexamos → tmux

```
root@drbd-01--
DRBD DrbdMon v9.12.2 | Node drbd-01.cadilinea.lan
RES: drbd0
drbd-02
Primary 0: 0
Page: 1
Quit Reprint Pg1 PgUp PgDn Problems
DRBD DrbdMon v9.12.2 | Node drbd-02.cadilinea.lan
RES: drbd0
drbd-01
Primary 0: 0
Page: 1
Quit Reprint Pg1 PgUp PgDn Problems
[0] 0:root@drbd-02:~* "LINBIT* DrbdMon(Node " 09:56 07-May-21
```

→ Instalación → drbdtop

[https://github.com/hdjr/install\\_drbdtop/blob/master/install\\_drbdtop.sh](https://github.com/hdjr/install_drbdtop/blob/master/install_drbdtop.sh)

# yum install git

# git clone [https://github.com/hdjr/install\\_drbdtop.git](https://github.com/hdjr/install_drbdtop.git)

# sh install\_drbdtop/install\_drbdtop.sh

# su -

# drbdtop → Multiplexamos → tmux

```

root@drbd-01~
DRBDTOP v0.2.3 (kernel: 9.0.22; utils: 9.12.2; host: drbd-01.cadilinea.lan)
(LIVE UPDATING) Resource List
┌───┬───┬───┬───┬───┬───┬───┐
│ Name │ Role │ Disks │ Peer Disks │ Connections │ Overall │ Quorum │
├───┬───┬───┬───┬───┬───┬───┤
│ drbd0 │ Primary │ ✓ │ ✓ │ ✓ │ ✓ │ ✓ │
└───┬───┬───┬───┬───┬───┬───┘

q: QUIT | j/k: down/up | f: Toggle dangerous filter | <tab>: Toggle updates

root@drbd-01~
DRBDTOP v0.2.3 (kernel: 9.0.22; utils: 9.12.2; host: drbd-02.cadilinea.lan)
(LIVE UPDATING) Resource List
┌───┬───┬───┬───┬───┬───┬───┐
│ Name │ Role │ Disks │ Peer Disks │ Connections │ Overall │ Quorum │
├───┬───┬───┬───┬───┬───┬───┤
│ drbd0 │ Primary │ ✓ │ ✓ │ ✓ │ ✓ │ ✓ │
└───┬───┬───┬───┬───┬───┬───┘

q: QUIT | j/k: down/up | f: Toggle dangerous filter | <tab>: Toggle updates
[0] 0:root@drbd-02:~* "root@drbd-01:~" 10:30 08-May-21

```

→ **DRBD** → **drbdmanage**

(**drbdmanage** es una alternativa mas simple de administración de recursos, peers, split-brain, ...).

```
# yum install wget pyobject2
```

```
# wget http://www.linbit.com/downloads/drbdmanage/drbdmanage-0.99.18.tar.gz
```

```
# tar -xvzf drbdmanage-0.99.18.tar.gz -C /opt/
```

```
# cd /opt/drbdmanage-0.99.18/
```

```
drbdmanage-0.99.18 # ./setup.py build
```

```
drbdmanage-0.99.18 # ./setup.py install
```

```
drbdmanage-0.99.18 # cd
```

```
# rm drbdmanage-0.99.18.tar.gz
```

```
# drbdmanage ping
```

```
pong
```

```
# su -
```

→ **Montaje Manual** → **xfs|FS**

```
# lsblk /dev/vdb
```

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

```
vdb           252:16  0  20G  0 disk
```

```
└─vdb1      252:17  0  20G  0 part
  └─vg_drbd-lv_drbd 253:5   0   1G  0 lvm
    └─drbd0    147:0   0 1024M  0 disk
```

```
# mkdir /mnt/drbd0 -p
```

```
# mkfs.xfs /dev/drbd0
```

```
meta-data=/dev/drbd0      isize=512  agcount=4, agsize=65532 blks
```

```
        =                sectsz=512  attr=2, projid32bit=1
```

```
        =                crc=1      finobt=0, sparse=0
```

```
data      =                bsize=4096  blocks=262127, imaxpct=25
```

```
        =                sunit=0     swidth=0 blks
```

```
naming    =version 2        bsize=4096  ascii-ci=0 ftype=1
```

```
log       =internal log    bsize=4096  blocks=855, version=2
```

```
        =                sectsz=512  sunit=0 blks, lazy-count=1
```

```
realtime  =none           extsz=4096  blocks=0, rtextents=0
```

```
# mount /dev/drbd0 /mnt/drbd0/
```

```
# df -hT /dev/drbd0
```

```
S.ficheros  Tipo Tamaño Usados  Disp Uso% Montado en
/dev/drbd0  xfs  1021M  33M  989M   4% /mnt/drbd0
```

```
# touch /mnt/drbd0/Prueba.txt
```

```
drbd-01 ~ # ls /mnt/drbd0/
```

```
Prueba.txt
```

```
drbd-02 ~ # ls /mnt/drbd0/
```

```
Prueba.txt
```

→ [Integración de DRBD con Pacemaker.](#)

```
# yum install pacemaker pcs
```

```
# firewall-cmd --permanent --add-service=high-availability
```

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

```
# systemctl enable --now pcsd.service
# passwd hacluster
drbd-01 ~ # pcs cluster auth drbd-01 drbd-02 -u hacluster
drbd-01 ~ # pcs cluster setup --name ha_drbd drbd-01 drbd-02
```

Destroying cluster on nodes: drbd-01, drbd-02...

drbd-01: Stopping Cluster (pacemaker)...

drbd-02: Stopping Cluster (pacemaker)...

drbd-02: Successfully destroyed cluster

drbd-01: Successfully destroyed cluster

Sending 'pacemaker\_remote authkey' to 'drbd-01', 'drbd-02'

drbd-01: successful distribution of the file 'pacemaker\_remote authkey'

drbd-02: successful distribution of the file 'pacemaker\_remote authkey'

Sending cluster config files to the nodes...

drbd-01: Succeeded

drbd-02: Succeeded

Synchronizing pcsd certificates on nodes drbd-01, drbd-02...

drbd-02: Success

drbd-01: Success

Restarting pcsd on the nodes in order to reload the certificates...

drbd-02: Success

drbd-01: Success

```
drbd-01 ~ # pcs cluster start --all
drbd-01 ~ # pcs cluster enable --all
drbd-01 ~ # pcs property set stonith-enabled=false
drbd-01 ~ # pcs status cluster
```

Cluster Status:

Stack: corosync

Current DC: drbd-01 (version 1.1.23-1.el7\_9.1-9acf116022) - partition with quorum

Last updated: Sat May 8 11:28:34 2021

Last change: Sat May 8 11:27:51 2021 by hacluster via crmd on drbd-01

2 nodes configured

0 resource instances configured

PCSD Status:

drbd-01: Online

drbd-02: Online

**drbd-01 ~ # pcs property**

Cluster Properties:

cluster-infrastructure: corosync

cluster-name: ha\_drbd

dc-version: 1.1.23-1.el7\_9.1-9acf116022

have-watchdog: false

stonith-enabled: false

→ [Administración del Cluster](#) → **pacemaker**

→ Los recursos y el FS|xfs deben ser administrados ahora por → **pacemaker**.

→ **Desmontamos, Paramos, y deshabilitamos para ceder el control** → **pacemaker**.

**# umount /mnt/drbd0**

**# systemctl disable drbd.service**

**# systemctl stop drbd.service**

**drbd-01 ~ # pcs property set no-quorum-policy=ignore**

**drbd-01 ~ # pcs property set default-resource-stickiness="INFINITY"**

**drbd-01 ~ # pcs property**

Cluster Properties:

cluster-infrastructure: corosync

cluster-name: **ha\_drbd**

dc-version: 1.1.23-1.el7\_9.1-9acf116022

default-resource-stickiness: INFINITY

have-watchdog: false

no-quorum-policy: ignore

stonith-enabled: false

**drbd-01 ~ # pcs resource create ClusterIP ocf:heartbeat:IPaddr2 ip=192.168.10.150  
cidr\_netmask=24 op monitor interval=30s**

**drbd-01 ~ # pcs resource describe ocf:linbit:drbd**

ocf:linbit:drbd - Manages a DRBD device as a Master/Slave resource

This resource agent manages a DRBD resource as a master/slave resource.

DRBD is a shared-nothing replicated storage device.

NOTE:

To avoid data-divergence, you should enable either

DRBD "quorum" and "on-no-quorum io-error" (recommended),

or configure proper fencing policies in both DRBD

\*and\* Pacemaker (fencing resource-and-stonith).

This cannot be done from this resource agent alone.

See the DRBD User's Guide for more information.

<https://docs.linbit.com/>

Resource options:

**drbd\_resource** (required) (unique): The name of the drbd resource from the drbd.conf file.

**drbdconf**: Full path to the drbd.conf file.

**adjust\_master\_score**: Space separated list of four master score adjustments for different scenarios:

- only access to 'consistent' data

- only remote access to 'uptodate' data - currently Secondary, local access to 'uptodate'

data, but remote is

unknown - local access to 'uptodate' data, and currently Primary or remote is known

Numeric values are expected

to be non-decreasing. The first value is 0 by default to prevent pacemaker from trying to promote while it is

unclear whether the data is really the most recent copy. (DRBD knows it is "consistent", but is unsure about

"uptodate"ness). Please configure proper fencing methods both in DRBD (fencing resource-and-stonith; appropriate

(un)fence-peer handlers) AND in Pacemaker to make this work reliably. Advanced use: Adjust the other values to

better fit into complex dependency score calculations. Intentionally diskless nodes ("Diskless Clients") with

access to good data via some (or all) their peers will use the 3rd or 4th value (minus one) when they are

(Secondary, not all peers up-to-date) or (ALL peers are up-to-date, or they are Primary themselves). This may

need to change if this should become a frequent use case. Special considerations: If a Secondary DRBD is

connected to a peer in Primary role, but Pacemaker does not know about any Primary (using `crm_resource --locate`),

we conclude that there likely is a cluster-split-brain, and may try to "help" Pacemaker by removing the master-

score. Also see "remove\_master\_score\_if\_peer\_primary".

`stop_outdates_secondary`: Recommended setting: leave at default (disabled). Note that this feature depends on the passed in information

in `OCF_RESKEY_CRM_meta_notify_master_uname` to be correct, which unfortunately is not reliable for pacemaker

versions up to at least 1.0.10 / 1.1.4. If a Secondary is stopped (unconfigured), it may be marked as

outdated in the drbd meta data, if we know there is still a Primary running in the cluster. Note that this

does not affect fencing policies set in drbd config, but is an additional safety feature of this resource

agent only. You can enable this behaviour by setting the parameter to true. If this feature seems to not do

what you expect, make sure you have defined fencing policies in the drbd configuration as well.

`ignore_missing_notifications`: Some setups do not benefit from notifications. Allow to disable notifications without patching this

resource agent.

`wfc_timeout`: Unless set to the empty string or any non-digits, wait (at most) this many seconds for the connection(s) to be

established after bringing them up during "start".

`remove_master_score_if_peer_primary`: See also "adjust\_master\_score" and "fail\_promote\_early\_if\_peer\_primary". To prevent a potentially

failed promotion attempt in case of cluster split-brain (Pacemaker communication loss) while DRBD

is still connected to a Primary, you can request to remove any master score while DRBD is

connected to a Primary (and that Primary peer looks like it has all disks up-to-date). This may

delay legitimate failovers after Primary crash by up to some TCP timeout (until DRBD realizes

that the Primary is gone) plus one monitoring interval. This parameter is interpreted almost as

an "ocf boolean", with the exception of a literal "unexpected", that is: - (yes|true|1)

[actually, according to the OCF spec, also (YES|TRUE|True|ja|ON), but please don't go there]: is

"true": remove (or never assign) master scores, if DRBD appears to see a (healthy) Primary -

"unexpected": assign master scores as described under "adjust\_master\_score", while removing it if



DRBD appears to see a (healthy) Primary that Pacemaker does not know about (as determined by

`crm_resource --locate`). - everything else is "false": ignore the peer role while assigning master scores.

`fail_promote_early_if_peer_primary`: See also "adjust\_master\_score" and "remove\_master\_score\_if\_peer\_primary". To avoid a useless retry

loop during promotion attempts in case of cluster split-brain (Pacemaker communication loss) while

DRBD is still connected to a Primary, you can chose to give up after the first try if this

situation is detected. If a Primary "vanishes", TCP may not immediately detect this, and an idle

DRBD may take some time until it does in-DRBD-protocol "pings". Pacemaker may well detect Primary

loss earlier than DRBD, and try to promote while DRBD thinks it can still see a Primary. Which

means, in general, trying to promote at least once is necessary, as that implies an in-DRBD-

protocol "peer alive" check. But if that does not succeed, re-trying until we hit the operation

timeout may not be desired, so you can disable it.

`unfence_if_all_uptodate`: If all volumes of this resource report to be UpToDate, call an unfence script hook, just in case some stale

fencing constraint or similar is still around. - With DRBD utils version  $\leq$  8.9.4, this is hardcoded to

`/usr/lib/drbd/crm-unfence-peer.sh -r $DRBD_RESOURCE` - With DRBD utils version  $\geq$  8.9.5, this is dispatched to

`$DRBDADM unfence-peer $DRBD_RESOURCE` In any case, the hook itself is responsible to fetch

`$OCF_RESKEY_unfence_extra_args` from its environment.

`unfence_extra_args`: This may be used to pass extra hints to the unfence hook. See description of

unfence\_if\_all\_uptodate.

`require_drbd_module_version_ge`: Use this you want to force failure of this resource agent if the detected DRBD kernel (module) driver

version is lower than a required minimum. Example: use

`require_drbd_module_version_ge=9.0.16` to fail

unless DRBD module version  $\geq 9.0.16$  is available (effectively requires DRBD 9). The intention of this

is to give a more useful failure message after accidentally downgrading the DRBD version by

installing/upgrading a new kernel. Note: "ge", "greater-or-equal", inclusive.

Required format: x.y.z

`require_drbd_module_version_lt`: Use this you want to force failure of this resource agent if the detected DRBD kernel (module) driver

version is higher than a required maximum. Example: use

`require_drbd_module_version_lt=9.0.0` to fail

unless DRBD module version  $< 9.0$  is available (effectively requires DRBD 8.4). Note: "lt", "less-

than", exclusive. Required format: x.y.z

`connect_only_after_promote`: This may be useful for "stacked" setups without proper fencing on the lower layer (which we obviously do

not recommend), to avoid some of the ugly side effects that may arise after resolving a split-brain on the

lower layer. Keep this DRBD instance disconnected until it is promoted. After promotion we issue an

additional "adjust", which is supposed to initiate the connection attempts. This causes a new data

generation identifier ("current uuid") to be generated after the failover of a "healthy" DRBD.

Default operations:

start: interval=0s timeout=240

reload: interval=0s timeout=30

promote: interval=0s timeout=90

```
demote: interval=0s timeout=90
notify: interval=0s timeout=90
stop: interval=0s timeout=100
monitor: interval=20 role=Slave timeout=20
monitor: interval=10 role=Master timeout=20
```

```
drbd-01 ~ # pcs resource create DrbdData ocf:linbit:drbd drbd_resource=drbd0 op monitor interval=60s
```

```
drbd-01 ~ # pcs resource master DrbdDataClone DrbdData master-max=1 master-node-max=1 clone-max=2 clone-node-max=1 notify=true
```

```
drbd-01 ~ # pcs resource create DrbdFS Filesystem device="/dev/drbd0" directory="/mnt/drbd0" fstype="xfs"
```

→ → **Posición de los recursos** → **LOC (location, order, colocation)**

→ **location**

```
drbd-01 ~ # pcs constraint location ClusterIP prefers drbd-01=INFINITY
```

```
drbd-01 ~ # pcs constraint location DrbdDataClone prefers drbd-01=INFINITY
```

```
drbd-01 ~ # pcs constraint location DrbdFS prefers drbd-01=INFINITY
```

→ **order**

```
drbd-01 ~ # pcs constraint order promote DrbdDataClone then start DrbdFS
```

→ **colocation**

```
drbd-01 ~ # pcs constraint colocation add DrbdFS with DrbdDataClone INFINITY with-rsc-role=Master
```

```
drbd-01 ~ # pcs constraint colocation add ClusterIP with DrbdFS INFINITY
```

```
drbd-01 ~ # pcs constraint
```

Location Constraints:

Resource: ClusterIP

Enabled on: drbd-01 (score:INFINITY)

Enabled on: drbd-01 (score:INFINITY) (role: Started)

Resource: DrbdDataClone

Enabled on: drbd-01 (score:INFINITY)



Last updated: Sat May 8 18:47:06 2021

Last change: Sat May 8 18:36:01 2021 by root via cibadmin on drbd-01

2 nodes configured

4 resource instances configured

Online: [ drbd-01 drbd-02 ]

Full list of resources:

ClusterIP (ocf::heartbeat:IPaddr2): Started drbd-01

Master/Slave Set: DrbdDataClone [DrbdData]

Masters: [ drbd-01 ]

Slaves: [ drbd-02 ]

DrbdFS (ocf::heartbeat:Filesystem): Started drbd-01

Daemon Status:

corosync: active/enabled

pacemaker: active/enabled

pcsd: active/enabled

#### **drbd-01 ~ # pcs resource**

ClusterIP (ocf::heartbeat:IPaddr2): Started drbd-01

Master/Slave Set: DrbdDataClone [DrbdData]

Masters: [ drbd-01 ]

Slaves: [ drbd-02 ]

DrbdFS (ocf::heartbeat:Filesystem): Started drbd-01

#### **drbd-01 ~ # pcs resource op defaults timeout=60s**

#### **drbd-01 ~ # pcs resource op defaults**

timeout=60s

**BIBLIOGRAFIA:**

<https://www.atlantic.net/vps-hosting/install-and-configure-drbd-on-centos-8/>

<https://www.learnitguide.net/2016/07/integrate-drbd-with-pacemaker-clusters.html>

<https://documentation.suse.com/sle-ha/12-SP4/html/SLE-HA-all/cha-ha-drbd.html>

<https://documentation.suse.com/sle-ha/12-SP4/html/SLE-HA-all/cha-ha-drbd.html>

**Creative Commons****Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA**

© 2021 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.