

**→ 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'.
```

**→ Recursos: → 'defaults' → pacemaker:**

```
[root@srv1 ~]# pcs resource defaults migration-threshold=5  
Warning: Defaults do not apply to resources which override them with their own defined values  
[root@srv1 ~]# pcs resource defaults resource-stickiness=10  
Warning: Defaults do not apply to resources which override them with their own defined values  
[root@srv2 ~]# pcs resource defaults failure-timeout=5  
Warning: Defaults do not apply to resources which override them with their own defined values
```

. **resource-stickiness**: adds a sticky score for the resource on its current node. It helps avoiding a resource move back and forth between nodes where it has the same score.

. **migration-threshold**: this controls how many time the cluster tries to recover a resource on the same node before moving it on another one.

. **failure-timeout** is measured since the most recent failure. That is, older failures do not individually time out and lower the fail count. Instead, all failures are timed out simultaneously (and the fail count is reset to 0) if there is no new failure for the timeout period.

**→ Ip → 'pgsql-vip' → Balanceada para: → 'team0':**

```
[root@srv1 ~]# pcs resource create pgsql-vip ocf:heartbeat:IPaddr2 nic="team0"  
ip=192.168.10.150 cidr_netmask=24 op monitor interval=30s
```

```
[root@srv1 ~]# pcs constraint location pgsql-vip prefers srv1.enermol.lan  
[root@srv1 ~]# pcs constraint location  
Location Constraints:  
Resource: pgsql-vip  
Enabled on:  
Node: srv1.enermol.lan (score:INFINITY)
```

**→ Desplegamos el clon → 'glusterd' → Cedemos el control de los montajes 'glusterFS' a 'pacemaker', para todos los servers que lo requieran y en modo 'cliente nativo' → 'fuse.glusterfs'.**

```
[root@srv1 ~]# pcs resource create glusterd systemd:glusterd  
[root@srv1 ~]# pcs resource clone glusterd
```

```
[root@srv1 ~]# pcs resource create glusterFS-srv1 ocf:heartbeat:Filesystem  
device="srv1.enermol.lan:/repVol-01" directory="/srv/repVol-01/" fstype="glusterfs"  
options="defaults"  
[root@srv1 ~]# pcs constraint location glusterFS-srv1 prefers srv1.enermol.lan=INFINITY
```

```
[root@srv1 ~]# pcs resource create glusterFS-srv2 ocf:heartbeat:Filesystem  
device="srv2.enermol.lan:/repVol-01" directory="/srv/repVol-01/" fstype="glusterfs"  
options="defaults"  
[root@srv1 ~]# pcs constraint location glusterFS-srv2 prefers srv2.enermol.lan=INFINITY
```



→ Desplegamos → 'postgresql' → El montaje de esta forma pretende evitar el diseño de un 'PAF' (postgres Active failover), diseñado ad-hoc para postgres, y en streaming puesto que las restricciones las evitará 'GlusterFS'. Este último será el responsable de evitar los split-brains. Confiamos por tanto en la selección de atributos extendidos en una capa inferior. Las BB. de DD. deben hacer su trabajo en otro nivel. No tiene ningún sentido buscar un testigo de réplica en una base de datos idénticamente replicada, puesto que el testigo sería siempre el mismo. Con lo cual la réplica no podría ser nunca fiable. Y provocaría errores de tokens sucesivos.

```
[root@srv1 ~]# pcs resource create postgresql systemd:postgresql
[root@srv1 ~]# pcs constraint location postgresql prefers srv1.enermol.lan
[root@srv1 ~]# pcs constraint colocation add pgsq1-vip with postgresql score=+INFINITY
[root@srv1 ~]# pcs constraint order set pgsq1-vip postgresql sequential=false require-all=true
```

→ Desplegamos → 'http' → apache. Con las restricciones impuestas se balanceará a favor de la 'pgsq1-vip'. De interés para un despliegue → 'pgadmin4'.

```
# vim /etc/httpd/conf.d/status.conf
<Location /server-status>
SetHandler server-status
Order deny,allow
Deny from all
Allow from 127.0.0.1
Allow from ::1
</Location>
```

```
[root@srv1 ~]# pcs resource create apache ocf:heartbeat:apache
configfile=/etc/httpd/conf/httpd.conf statusurl="http://localhost/server-status" op monitor
interval=1min
```

```
[root@srv1 ~]# pcs constraint location apache prefers srv1.enermol.lan
[root@srv1 ~]# pcs constraint order pgsq1-vip then apache
[root@srv1 ~]# pcs constraint colocation add apache with pgsq1-vip score=+INFINITY
```

→ Configuración Actual del Cluster → 'cluster-odoo'

```
# pcs config
Cluster Name: cluster-odoo
Corosync Nodes:
  srv1.enermol.lan srv2.enermol.lan
Pacemaker Nodes:
  srv1.enermol.lan srv2.enermol.lan
```

```
Resources:
Resource: pgsq1-vip (class=ocf provider=heartbeat type=IPAddr2)
Attributes: cidr_netmask=24 ip=192.168.10.150 nic=enp1s0
Operations: monitor interval=30s (pgsq1-vip-monitor-interval-30s)
```



start interval=0s timeout=20s (pgsql-vip-start-interval-0s)

stop interval=0s timeout=20s (pgsql-vip-stop-interval-0s)

Clone: glusterd-clone

Resource: glusterd (class=ocf provider=glusterfs type=glusterd)

Operations: monitor interval=10 timeout=20 (glusterd-monitor-interval-10)

reload interval=0s timeout=20 (glusterd-reload-interval-0s)

start interval=0s timeout=20 (glusterd-start-interval-0s)

stop interval=0s timeout=20 (glusterd-stop-interval-0s)

Resource: glusterFS-srv1 (class=ocf provider=heartbeat type=Filesystem)

Attributes: device=srv1:/repVol-01 directory=/srv/repVol-01/ fstype=glusterfs

Operations: monitor interval=20s timeout=40s (glusterFS-srv1-monitor-interval-20s)

start interval=0s timeout=60s (glusterFS-srv1-start-interval-0s)

stop interval=0s timeout=60s (glusterFS-srv1-stop-interval-0s)

Resource: glusterFS-srv2 (class=ocf provider=heartbeat type=Filesystem)

Attributes: device=srv2:/repVol-01 directory=/srv/repVol-01/ fstype=glusterfs

Operations: monitor interval=20s timeout=40s (glusterFS-srv2-monitor-interval-20s)

start interval=0s timeout=60s (glusterFS-srv2-start-interval-0s)

stop interval=0s timeout=60s (glusterFS-srv2-stop-interval-0s)

Resource: postgresql (class=systemd type=postgresql)

Operations: monitor interval=60 timeout=100 (postgresql-monitor-interval-60)

start interval=0s timeout=100 (postgresql-start-interval-0s)

stop interval=0s timeout=100 (postgresql-stop-interval-0s)

Stonith Devices:

Fencing Levels:

Location Constraints:

Resource: glusterFS-srv1

Enabled on:

Node: srv1.enermol.lan (score:INFINITY) (id:location-glusterFS-srv1-srv1.enermol.lan-INFINITY)

Resource: glusterFS-srv2

Enabled on:

Node: srv2.enermol.lan (score:INFINITY) (id:location-glusterFS-srv2-srv2.enermol.lan-INFINITY)

Resource: pgsql-vip

Enabled on:

Node: srv1.enermol.lan (score:INFINITY) (id:location-pgsql-vip-srv1.enermol.lan-INFINITY)

Resource: postgresql

Enabled on:

Node: srv1.enermol.lan (score:INFINITY) (id:location-postgresql-srv1.enermol.lan-INFINITY)

Ordering Constraints:

Resource Sets:

set pgsql-vip postgresql require-all=true sequential=false (id:pcs\_rsc\_set\_pgsql-vip\_postgresql)  
(id:pcs\_rsc\_order\_set\_pgsql-vip\_postgresql)

Colocation Constraints:

pgsql-vip with postgresql (score:+INFINITY) (id:colocation-pgsql-vip-postgresql-INFINITY)



Ticket Constraints:

Alerts:

No alerts defined

### Resources Defaults:

[failure-timeout=5](#)

[migration-threshold=5](#)

[resource-stickiness=10](#)

Operations Defaults:

No defaults set

Cluster Properties:

cluster-infrastructure: corosync

cluster-name: cluster-odoo

dc-version: 2.0.3-5.el8\_2.1-4b1f869f0f

have-watchdog: false

last-lrm-refresh: 1597078285

stonith-enabled: false

**symmetric-cluster: true**

Quorum:

Options:

```
# pcs status
```

```
Cluster name: cluster-odoo
```

```
Cluster Summary:
```

```
* Stack: corosync
```

```
* Current DC: srv2.enermol.lan (version 2.0.3-5.el8_2.1-4b1f869f0f) -  
partition with quorum
```

```
* Last updated: Mon Aug 10 23:17:05 2020
```

```
* Last change: Mon Aug 10 18:51:25 2020 by hacluster via crmd on
```

```
srv2.enermol.lan
```

```
* 2 nodes configured
```

```
* 6 resource instances configured
```

```
Node List:
```

```
* Online: [ srv1.enermol.lan srv2.enermol.lan ]
```

```
Full List of Resources:
```

```
* pgsql-vip (ocf::heartbeat:IPaddr2): Started srv1.enermol.lan
```

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

```
* Started: [ srv1.enermol.lan srv2.enermol.lan ]
```

```
* glusterFS-srv1 (ocf::heartbeat:Filesystem): Started srv1.enermol.lan
```

```
* glusterFS-srv2 (ocf::heartbeat:Filesystem): Started srv2.enermol.lan
```

```
* postgresql (systemd:postgresql): Started srv1.enermol.lan
```

```
Daemon Status:
```

```
corosync: active/enabled
```

```
pacemaker: active/enabled
```

```
pcsd: active/enabled
```



## → Configuración Actual del Volúmen y bricks → 'repVol-01'

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

```
Volume Name: repVol-01
Type: Replicate
Volume ID: 84c89f93-9904-4aa2-8c48-3e8e7c3c3338
Status: Started
Snapshot Count: 0
Number of Bricks: 1 x 3 = 3
Transport-type: tcp
Bricks:
Brick1: srv1.enermol.lan:/bricks/repBrick-01/data
Brick2: srv2.enermol.lan:/bricks/repBrick-01/data
Brick3: srv3.enermol.lan:/bricks/repBrick-01/data
Options Reconfigured:
performance.client-io-threads: off
nfs.disable: on
storage.fips-mode-rchecksum: on
transport.address-family: inet
features.barrier-timeout: 2m
features.uss: enable
cluster.self-heal-daemon: on
cluster.entry-self-heal: on
cluster.metadata-self-heal: on
cluster.data-self-heal: on
cluster.quorum-type: auto
cluster.server-quorum-type: server
auth.allow: 192.168.10.*
features.bitrot: on
features.scrub: Active
features.barrier: enable
cluster.server-quorum-ratio: 51%
cluster.enable-shared-storage: enable
auto-delete: enable
```

**REFERENCIAS:****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:

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.