

→ 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'.
```

361.3 Failover Clusters (weight: 8)

Weight	8
Description	Candidates should have experience in the installation, configuration, maintenance and troubleshooting of a Pacemaker cluster. This includes the use of Corosync. The focus is on Pacemaker 2.x for Corosync 2.x.

Key Knowledge Areas:

- Understand the architecture and components of Pacemaker (CIB, CRMD, PEngine, LRMd, DC, STONITHd)
- Manage Pacemaker cluster configurations
- Understand Pacemaker resource classes (OCF, LSB, Systemd, Service, STONITH, Nagios)
- Manage Pacemaker resources
- Manage resource rules and constraints (location, order, colocation).
- Manage advanced resource features (templates, groups, clone resources, multi-state resources)
- Obtain node information and manage node health
- Manage quorum and fencing in a Pacemaker cluster
- Configure the Split Brain Detector on shared storage
- Manage Pacemaker using pcs
- Manage Pacemaker using crmsh
- Configure and management of corosync in conjunction with Pacemaker
- Awareness of Pacemaker ACLs
- Awareness of other cluster engines (OpenAIS, Heartbeat, CMAN)

Partial list of the used files, terms and utilities:

- pcs
- crm
- crm_mon
- crm_verify
- crm_simulate
- crm_shadow
- crm_resource
- crm_attribute
- crm_node
- crm_standby

- cibadmin
- corosync.conf
- authkey
- corosync-cfgtool
- corosync-cmapctl
- corosync-quorumtool
- stonith_admin
- stonith
- ocf:pacemaker:ping
- ocf:pacemaker:NodeUtilization
- ocf:pacemaker:ocf:SysInfo
- ocf:pacemaker:HealthCPU
- ocf:pacemaker:HealthSMART
- sbd

→ [pacemaker -> hosts.](#)

pcmk-XX ~ # vim /etc/hosts

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

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

```
## pacemaker -> hosts.
```

```
192.168.10.160 pcmk-00.cadilinea.lan pcmk-00
```

```
192.168.10.161 pcmk-01.cadilinea.lan pcmk-01
```

```
192.168.10.162 pcmk-02.cadilinea.lan pcmk-02
```

```
192.168.10.163 pcmk-03.cadilinea.lan pcmk-03
```

```
192.168.10.165 vip.cadilinea.lan vip
```

→ [Sincronización horaria.](#)

pcmk-XX ~ # dnf install chrony

pcmk-XX ~ # vim /etc/chrony.conf

...

```
# Allow NTP client access from local network.
```

```
allow 192.168.10.0/24
```

pcmk-XX ~ # timedatectl set-timezone Europe/Madrid

pcmk-XX ~ # systemctl restart chronyd

pcmk-XX ~ # systemctl status chronyd

pcmk-XX ~ # firewall-cmd --permanent --add-service=ntp

pcmk-XX ~ # firewall-cmd --permanent --add-port=123/udp

pcmk-XX ~ # firewall-cmd --reload

pcmk-XX ~ # chronyc tracking

Reference ID : 2561C3C3 (server.pasynike.nl)

Stratum : 3

Ref time (UTC) : Thu Apr 01 07:24:39 2021

System time : 0.000092607 seconds slow of NTP time

Last offset : +0.001440001 seconds

RMS offset : 0.001440001 seconds

Frequency : 16.384 ppm slow

Residual freq : -1.726 ppm

Skew : 39.363 ppm

Root delay : 0.080205671 seconds

Root dispersion : 0.032234624 seconds

Update interval : 64.3 seconds

Leap status : Normal

pcmk-XX ~ # chronyc sources

210 Number of sources = 4

MS Name/IP address	Stratum	Poll	Reach	LastRx	Last sample
=====					
^+ ns3.turbodns.co.uk	2	6	377	5	-1505us[-1505us] +/- 52ms
^+ 172.86.181.77	2	6	373	6	+6111us[+6111us] +/- 105ms
^- server.bertold.org	2	6	17	2	+862us[+862us] +/- 68ms
^* server.pasynike.nl	2	6	377	7	+4957us[+6397us] +/- 70ms

=====

pcmk-XX ~ # chronyc sourcestats

210 Number of sources = 4

Name/IP Address	NP	NR	Span	Frequency	Freq Skew	Offset	Std Dev
ns3.turbodns.co.uk	9	6	330	+9.508	23.570	+407us	1863us
172.86.181.77	8	4	330	+15.732	55.045	+6972us	2932us
server.bertold.org	5	3	257	-4.006	80.706	-89us	838us
server.pasynike.nl	9	5	330	+14.272	52.219	+3227us	3106us

=====

pcmk-XX ~ # chronyc makestep

200 OK

pcmk-XX ~ # timedatectl

Local time: jue 2021-04-01 09:27:49 CEST

Universal time: jue 2021-04-01 07:27:49 UTC

RTC time: jue 2021-04-01 07:27:49

Time zone: Europe/Madrid (CEST, +0200)

System clock synchronized: no

NTP service: active

RTC in local TZ: no

→ [Arquitectura pacemaker/corosync.](#)

→ **CIB** (Cluster Information Base). Es un daemon XML. La configuración la utilizará el **DC** (Designed Coordinator).

→ **CRMd** (cluster Resource Management daemon). Dirige las acciones del cluster. Existe también **LRMd** local que interactúa con los agentes CRMd.

→ **STONITH** (Shoot The Other Node In The Head). Propósito de cercado o fencing.

→ **corosync**. Componente y daemon con el mismo nombre. Tareas:

- . Comunicación de componentes.

- . Quorum.

- . Capacidad de mensajes.

- . **kronosnet** → Librería de transporte de red para proporcionar múltiple redundancia.

→ [Instalación pacemaker/corosync](#).

```
pcm-XX ~ # /etc/ssh/sshd_config
```

...

```
PermitRootLogin yes
```

```
PasswordAuthentication yes
```

...

```
pcm-XX ~ # dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

```
pcm-XX ~ # dnf update
```

```
pcm-XX ~ # subscription-manager repos --enable=rhel-8-for-x86_64-highavailability-rpms
```

El repositorio 'rhel-8-for-x86_64-highavailability-rpms' está habilitado para este sistema.

```
pcm-XX ~ # dnf update
```

```
pcm-XX ~ # dnf repolist
```

Updating Subscription Management repositories.

id del repositorio	nombre del repositorio
--------------------	------------------------

epel	Extra Packages for Enterprise Linux 8 - x86_64
------	--

epel-modular	Extra Packages for Enterprise Linux Modular 8 - x86_64
--------------	--

rhel-8-for-x86_64-appstream-rpms	Red Hat Enterprise Linux 8 for x86_64 - AppStream (RPMs)
----------------------------------	--

rhel-8-for-x86_64-baseos-rpms	Red Hat Enterprise Linux 8 for x86_64 - BaseOS (RPMs)
-------------------------------	---

rhel-8-for-x86_64-highavailability-rpms	Red Hat Enterprise Linux 8 for x86_64 - High Availability (RPMs)
---	--

```
pcm-XX ~ # firewall-cmd --permanent --add-service=high-availability
```

```
pcm-XX ~ # firewall-cmd --reload
```

```
pcm-XX ~ # dnf install pacemaker pcs corosync
```

```
pcm-XX ~ # dnf install fence-agents-all pcp-zeroconf pcp-system-tools
```

→ **PCP – ZEROCONFIG** (Performance CoPilot). Herramienta de monitorización RedHat.

```
pcm-XX ~ # systemctl enable --now pcsd.service
```

```
pcm-XX ~ # passwd hacluster
```

→ [Configuración pacemaker/corosync](#).

```
pcm00 ~ # ssh-keygen
```

```
pcm00 ~ # ssh-copy-id -i .ssh/id_rsa.pub root@pcm01
```

```
pcm00 ~ # pcs host auth pcm00 pcm01 -u hacluster
```

Password:

```
pcm01: Authorized
```

```
pcm00: Authorized
```

```
pcm00 ~ # pcs cluster setup cluster00 pcm00 pcm01
```

```
No addresses specified for host 'pcm00', using 'pcm00'
```

```
No addresses specified for host 'pcm01', using 'pcm01'
```

```
Destroying cluster on hosts: 'pcm00', 'pcm01'...
```

```
pcm01: Successfully destroyed cluster
```

```
pcm00: Successfully destroyed cluster
```

```
Requesting remove 'pcsd settings' from 'pcm00', 'pcm01'
```

```
pcm00: successful removal of the file 'pcsd settings'
```

```
pcm01: successful removal of the file 'pcsd settings'
```

```
Sending 'corosync authkey', 'pacemaker authkey' to 'pcm00', 'pcm01'
```

```
pcm00: successful distribution of the file 'corosync authkey'
```

```
pcm00: successful distribution of the file 'pacemaker authkey'
```

```
pcm01: successful distribution of the file 'corosync authkey'
```

```
pcm01: successful distribution of the file 'pacemaker authkey'
```

```
Sending 'corosync.conf' to 'pcm00', 'pcm01'
```

```
pcm00: successful distribution of the file 'corosync.conf'
```

```
pcm01: successful distribution of the file 'corosync.conf'
```

```
Cluster has been successfully set up.
```

```
pcm00 ~ # vim /etc/corosync/corosync.conf
```

```
totem {
```

```
version: 2
cluster_name: cluster-00
transport: knet
crypto_cipher: aes256
crypto_hash: sha256
}

nodelist {
  node {
    ring0_addr: pcmk-00
    name: pcmk-00
    nodeid: 1
  }
  node {
    ring0_addr: pcmk-01
    name: pcmk-01
    nodeid: 2
  }
}

quorum {
  provider: corosync_votequorum
  two_node: 1
}

logging {
  to_logfile: yes
  logfile: /var/log/cluster/corosync.log
  to_syslog: yes
  timestamp: on
```

```
}
```

pcmk-00 ~ # corosync-keygen

Corosync Cluster Engine Authentication key generator.

Gathering 2048 bits for key from /dev/urandom.

Writing corosync key to /etc/corosync/authkey.

pcmk-00 ~ # pcs cluster start --all

pcmk-01: Starting Cluster...

pcmk-00: Starting Cluster...

pcmk-00 ~ # pcs cluster enable --all

pcmk-00: Cluster Enabled

pcmk-01: Cluster Enabled

pcmk-00 ~ # pcs status

Cluster name: cluster-00

WARNINGS:

No stonith devices and stonith-enabled is not false

Cluster Summary:

- * Stack: corosync

- * Current DC: pcmk-00 (version 2.0.4-6.el8_3.1-2deceaa3ae) - partition with quorum

- * Last updated: Thu Apr 1 11:19:21 2021

- * Last change: Thu Apr 1 11:18:25 2021 by hacluster via crmd on pcmk-00

- * 2 nodes configured

- * 0 resource instances configured

Node List:

- * Online: [pcmk-00 pcmk-01]

Full List of Resources:

- * No resources

Daemon Status:

corosync: active/enabled

pacemaker: active/enabled

pcsd: active/enabled

→ **Configuración XML/backups.**

pcmk-00 ~ # pcs cluster cib configuracion.xml

pcmk-00 ~ # cat configuracion.xml

```
<cib crm_feature_set="3.4.1" validate-with="pacemaker-3.4" epoch="5" num_updates="4"
admin_epoch="0" cib-last-written="Thu Apr 1 11:18:25 2021" update-origin="pcmk-00" update-
client="crmd" update-user="hacluster" have-quorum="1" dc-uuid="1">
```

```
<configuration>
```

```
<crm_config>
```

```
<cluster_property_set id="cib-bootstrap-options">
```

```
<nvpair id="cib-bootstrap-options-have-watchdog" name="have-watchdog" value="false"/>
```

```
<nvpair id="cib-bootstrap-options-dc-version" name="dc-version" value="2.0.4-6.el8_3.1-
2decea3ae"/>
```

```
<nvpair id="cib-bootstrap-options-cluster-infrastructure" name="cluster-infrastructure"
value="corosync"/>
```

```
<nvpair id="cib-bootstrap-options-cluster-name" name="cluster-name" value="cluster-00"/>
```

```
</cluster_property_set>
```

```
</crm_config>
```

```
<nodes>
```

```
<node id="1" uname="pcmk-00"/>
```

```
<node id="2" uname="pcmk-01"/>
```

...

pcmk-00 ~ # pcs cluster cib <TAB>

cib cib-push cib-upgrade

pcmk-00 ~ # pcs config <TAB>

backup checkpoint export import-cman restore

pcmk-00 ~ # pcs config backup configuración_cluster-00

pcmk-00 ~ # ls

anaconda-ks.cfg configuración_cluster-00.tar.bz2 configuracion.xml

→ **cibadmin** - query and edit the Pacemaker configuration

pcmk-00 ~ # cibadmin --query > tmp.xml

vi tmp.xml

cibadmin --replace --xml-file tmp.xml

cibadmin --query --scope resources > tmp.xml

vi tmp.xml

cibadmin --replace --scope resources --xml-file tmp.xml

cibadmin --query | grep stonith

```
<nvpair id="cib-bootstrap-options-stonith-action" name="stonith-action" value="reboot"/>
<nvpair id="cib-bootstrap-options-stonith-enabled" name="stonith-enabled" value="1"/>
<primitive id="child_DoFencing" class="stonith" type="external/vmware">
<lrms_resource id="child_DoFencing:0" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:0" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:1" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:0" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:2" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:0" type="external/vmware" class="stonith">
<lrms_resource id="child_DoFencing:3" type="external/vmware" class="stonith">
```

cibadmin --delete --xml-text '<primitive id="child_DoFencing"/>'

Ejemplos:

Query the configuration from the local node:

```
# cibadmin --query --local
```

Query just the cluster options configuration:

```
# cibadmin --query --scope crm_config
```

Query all 'target-role' settings:

```
# cibadmin --query --xpath "//nvpair[@name='target-role']"
```

Remove all 'is-managed' settings:

```
# cibadmin --delete-all --xpath "//nvpair[@name='is-managed']"
```

Remove the resource named 'old':

```
# cibadmin --delete --xml-text '<primitive id="old"/>'
```

Remove all resources from the configuration:

```
# cibadmin --replace --scope resources --xml-text '<resources/>'
```

Replace complete configuration with contents of \$HOME/pacemaker.xml:

```
# cibadmin --replace --xml-file $HOME/pacemaker.xml
```

Replace constraints section of configuration with contents of \$HOME/constraints.xml:

```
# cibadmin --replace --scope constraints --xml-file $HOME/constraints.xml
```

Increase configuration version to prevent old configurations from being loaded accidentally:

```
# cibadmin --modify --xml-text '<cib admin_epoch="admin_epoch++"/>'
```

Edit the configuration with your favorite \$EDITOR:

```
# cibadmin --query > $HOME/local.xml
```

```
# $EDITOR $HOME/local.xml
```

```
# cibadmin --replace --xml-file $HOME/local.xml
```

→ **crmadmin** - query and manage the Pacemaker controller

```
pcmk-00 ~ # crm <TAB>
```

```
crmadmin    crm_error    crm_mon      crm_resource  crm_simulate  crm_verify
```

```
crm_attribute  crm_failcount  crm_node     crm_rule      crm_standby
```

```
crm_diff      crm_master     crm_report   crm_shadow    crm_ticket
```

```
pcmk-00 ~ # crm_mon
```

Cluster Summary:

* Stack: corosync

* Current DC: pcmk-00 (version 2.0.4-6.el8_3.1-2deceaa3ae) - partition with quorum

* Last updated: Tue May 25 17:46:40 2021

* Last change: Tue May 25 17:21:17 2021 by hacluster via crmd on pcmk-00

* 2 nodes configured

* 0 resource instances configured

Node List:

* Online: [pcmk-00 pcmk-01]

Active Resources:

* No active resources

<CTRL> C

pcmk-00 ~ # crm_verify -LV

(unpack_resources) error: Resource start-up disabled since no STONITH resources have been defined

(unpack_resources) error: Either configure some or disable STONITH with the stonith-enabled option

(unpack_resources) error: NOTE: Clusters with shared data need STONITH to ensure data integrity

Errors found during check: config not valid

pcmk-00 ~ # crm_simulate --help

Usage:

crm_simulate [OPTION?]

crm_simulate - simulate a Pacemaker cluster's response to events

pcmk-00 ~ # crm_shadow

crm_shadow - perform Pacemaker configuration changes in a sandbox

pcmk-00 ~ # crm_shadow --create myShadow

Setting up shadow instance

Type Ctrl-D to exit the crm_shadow shell

```
shadow[myShadow] #
```

pcmk-00 ~ # crm_resource --help

Examples:

List the available OCF agents:

```
# crm_resource --list-agents ocf
```

List the available OCF agents from the linux-ha project:

```
# crm_resource --list-agents ocf:heartbeat
```

Move 'myResource' to a specific node:

```
# crm_resource --resource myResource --move --node altNode
```

Allow (but not force) 'myResource' to move back to its original location:

```
# crm_resource --resource myResource --clear
```

Stop 'myResource' (and anything that depends on it):

```
# crm_resource --resource myResource --set-parameter target-role --meta --parameter-value Stopped
```

Tell the cluster not to manage 'myResource' (the cluster will not attempt to start or stop the resource under any circumstances; useful when performing maintenance tasks on a resource):

```
# crm_resource --resource myResource --set-parameter is-managed --meta --parameter-value false
```

Erase the operation history of 'myResource' on 'aNode' (the cluster will 'forget' the existing resource state, including any errors, and attempt to recover the resource; useful when a resource had failed permanently and has been repaired by an administrator):

```
# crm_resource --resource myResource --cleanup --node aNode
```

pcmk-00 ~ # crmadmin

crmadmin - query and manage the Pacemaker controller

pcmk-00 ~ # crmadmin -N

member node: pcmk-00 (1)

member node: pcmk-01 (2)

pcmk-00 ~ # crm_attribute

crm_attribute - query and update Pacemaker cluster options and node attributes

pcmk-00 ~ # crm_attribute --node pcmk-00 --name location --query

scope=nodes name=location value=(null)

pcmk-00 ~ # crm_node

Usage:

```
crm_node [OPTION...]
```

crm_node - Tool for displaying low-level node information

Help Options:

- h, --help Show help options
- help-all Show all help options
- help-commands Show command help
- help-additional Show additional options

Application Options:

- \$, --version Display software version and exit
- V, --verbose Increase debug output (may be specified multiple times)
- Q, --quiet Be less descriptive in output.

Report bugs to users@clusterlabs.org

pcmk-00 ~ # corosync-cmapctl

```
config.totemconfig_reload_in_progress (u8) = 0
internal_configuration.service.0.name (str) = corosync_cmap
internal_configuration.service.0.ver (u32) = 0
internal_configuration.service.1.name (str) = corosync_cfg
internal_configuration.service.1.ver (u32) = 0
internal_configuration.service.2.name (str) = corosync_cpg
internal_configuration.service.2.ver (u32) = 0
internal_configuration.service.3.name (str) = corosync_quorum
internal_configuration.service.3.ver (u32) = 0
internal_configuration.service.4.name (str) = corosync_pload
```

pcmk-00 ~ # corosync-cfgtool -a pcmk-00

```
192.168.10.160
```

pcmk-00 ~ # corosync-quorumtool

```
Quorum information
```

```
-----
```

Date: Tue May 25 18:26:07 2021
Quorum provider: corosync_votequorum
Nodes: 2
Node ID: 1
Ring ID: 1.28
Quorate: Yes

Votequorum information

Expected votes: 2
Highest expected: 2
Total votes: 2
Quorum: 1
Flags: 2Node Quorate WaitForAll

Membership information

Nodeid Votes Name
1 1 pcmk-00 (local)
2 1 pcmk-01

pcmk-00 ~ # stonith_admin

Usage:

stonith_admin [OPTION...]

stonith_admin - Access the Pacemaker fencing API

pcmk-00 ~ # man sbd

NAME

sbd - STONITH Block Device daemon

SYNOPSIS

```
sbd <-d /dev/...> [options] "command"
```

SUMMARY

SBD provides a node fencing mechanism (Shoot the other node in the head, STONITH) for Pacemaker-based

clusters through the exchange of messages via shared block storage such as for example a SAN, iSCSI,

FCoE. This isolates the fencing mechanism from changes in firmware version or dependencies on specific

firmware controllers, and it can be used as a STONITH mechanism in all configurations that have

reliable shared storage.

pcmk-00 ~ # crm_standby

```
scope=status name=standby value=off
```

BIBLIOGRAFIA:

<https://www.itzgeek.com/post/how-to-setup-high-availability-cluster-on-centos-8-rhel-8/>

https://www.server-world.info/en/note?os=CentOS_8&p=pacemaker&f=1

https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/2.0/html/Pacemaker_Administration/s-cibadmin.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.