

**→ 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'.
```

→ Configuración 'Vagrant' inicial como ejemplo, no obstante utilizaré una configuración mas amplia y compleja para poder desplegar datos/servicios en entornos de trabajo diferenciados para un adecuado despliegue y en producción → Utilizaré KVM-queemu. Que sirva de referencia.

```
$ cat <<FIN > Vagrantfile  
Vagrant.configure(2) do |config|  
  config.vm.define "srv1" do |srv|  
    srv.vm.box = "centos/8"  
    srv.vm.hostname = "srv1.enermol.lan"  
    srv.vm.network "private_network", ip: "192.168.10.151",  
      virtualbox__intnet: "intnet"  
    srv.vm.network "private_network", ip: "192.168.10.161",  
      virtualbox__intnet: "intnet2"  
    srv.vm.network "forwarded_port", guest: 2224, host: 2224  
  end  
  config.vm.define "srv2" do |srv|  
    srv.vm.box = "centos/8"  
    srv.vm.hostname = "srv2.enermol.lan"  
    srv.vm.network "private_network", ip: "192.168.10.152",  
      virtualbox__intnet: "intnet"  
    srv.vm.network "private_network", ip: "192.168.10.162",  
      virtualbox__intnet: "intnet2"  
  end  
end  
FIN
```

```
# lsblk  
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT  
sda          8:0    0   40G  0 disk  
├─sda1       8:1    0    1G  0 part /boot  
└─sda2       8:2    0   39G  0 part  
   └─cl-root 253:0   0   36G  0 lvm  /  
     └─cl-swap 253:1   0  3,1G  0 lvm  [SWAP]  
sdb          8:16   0   40G  0 disk
```

```
# df -hT  
S.ficheros      Tipo      Tamaño Usados  Disp  Uso% Montado en  
devtmpfs        devtmpfs  1,4G    0      1,4G   0% /dev  
tmpfs           tmpfs     1,4G    38M    1,4G   3% /dev/shm  
tmpfs           tmpfs     1,4G    8,6M   1,4G   1% /run  
tmpfs           tmpfs     1,4G    0      1,4G   0% /sys/fs/cgroup  
/dev/mapper/cl-root xfs       36G    2,8G   34G    8% /  
/dev/sda1       ext4      976M   190M   720M   21% /boot  
tmpfs           tmpfs     284M    0      284M   0% /run/user/0
```

```
# cat <<FIN > /etc/hosts  
#127.0.0.1 srv1 srv1  
#127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4  
#::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
```



```
192.168.10.150 pgsq1-vip.enermol.lan pgsq1-vip
192.168.10.160 pgsq1-alt-vip.enermol.lan pgsq1-alt-vip
```

```
192.168.10.151 srv1.enermol.lan srv1
192.168.10.152 srv2.enermol.lan srv2
```

```
192.168.10.161 srv1-alt.enermol.lan srv1-alt
192.168.10.162 srv2-alt.enermol.lan srv2-alt
```

```
192.168.10.13 ilo-srv1.enermol.lan ilo-srv1
192.168.10.14 ilo-srv2.enermol.lan ilo-srv2
```

FIN

```
# vim /etc/ssh/sshd_config
```

...

```
PermitRootLogin yes
```

```
PasswordAuthentication yes
```

...

```
# systemctl restart sshd.service
```

```
# ssh-keygen
```

```
# systemctl status firewalld.service
```

```
● firewalld.service - firewalld - dynamic firewall daemon
  Loaded: loaded (/usr/lib/systemd/system/firewalld.service; enabled; vendor
 preset: enabled)
  Active: active (running) since Tue 2020-03-24 11:37:28 CET; 3h 52min ago
    Docs: man:firewalld(1)
  Main PID: 942 (firewalld)
    Tasks: 3 (limit: 17950)
   Memory: 41.4M
   CGroup: /system.slice/firewalld.service
           └─942 /usr/libexec/platform-python -s /usr/sbin/firewalld --nofork -
             nopic
```

```
mar 24 11:37:16 srv1.enermol.lan systemd[1]: Starting firewalld - dynamic
 firewall daemon...
```

```
mar 24 11:37:28 srv1.enermol.lan systemd[1]: Started firewalld - dynamic
 firewall daemon.
```

```
[root@srv1 ~]# ssh-copy-id -i .ssh/id_rsa.pub srv2
```

```
[root@srv2 ~]# ssh-copy-id -i .ssh/id_rsa.pub srv1
```

```
# 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
```

```
#
```

```
# systemctl restart chronyd.service chrony-wait.service
```

```
# chronyc tracking
```

```
# chronyc sources
```

```
# timedatectl set-timezone Europe/Madrid
```



```
# chronyc sourcestats
```

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/system_administrators_guide/sect-using_chrony

<http://proyectoa.com/configurar-actualizacion-automatica-de-hora-y-fecha-en-equipo-linux-centos-7-con-ntp/>

→ [Instalación de Pacemaker.](#)

```
# dnf install epel-release-8-8.el8.noarch -y
# firewall-cmd --permanent --add-service=high-availability
# firewall-cmd --reload
# firewall-cmd --list-services
cockpit dhcpv6-client high-availability ntp ssh

# dnf --enablerepo=HighAvailability -y install pacemaker pcs -y

# systemctl enable --now pcsd.service

# passwd hacluster
```

```
[root@srv1 ~]# pcs host auth srv1.enermol.lan srv2.enermol.lan -u hacluster
Password:
srv2.enermol.lan: Authorized
srv1.enermol.lan: Authorized
```

```
[root@srv1 ~]# pcs cluster setup cluster-odoo srv1.enermol.lan srv2.enermol.lan
No addresses specified for host 'srv1.enermol.lan', using 'srv1.enermol.lan'
No addresses specified for host 'srv2.enermol.lan', using 'srv2.enermol.lan'
Destroying cluster on hosts: 'srv1.enermol.lan', 'srv2.enermol.lan'...
srv1.enermol.lan: Successfully destroyed cluster
srv2.enermol.lan: Successfully destroyed cluster
Requesting remove 'pcsd settings' from 'srv1.enermol.lan', 'srv2.enermol.lan'
srv1.enermol.lan: successful removal of the file 'pcsd settings'
srv2.enermol.lan: successful removal of the file 'pcsd settings'
Sending 'corosync authkey', 'pacemaker authkey' to 'srv1.enermol.lan',
'srv2.enermol.lan'
srv1.enermol.lan: successful distribution of the file 'corosync authkey'
srv1.enermol.lan: successful distribution of the file 'pacemaker authkey'
srv2.enermol.lan: successful distribution of the file 'corosync authkey'
srv2.enermol.lan: successful distribution of the file 'pacemaker authkey'
Sending 'corosync.conf' to 'srv1.enermol.lan', 'srv2.enermol.lan'
srv2.enermol.lan: successful distribution of the file 'corosync.conf'
srv1.enermol.lan: successful distribution of the file 'corosync.conf'
Cluster has been successfully set up.
```

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

```
totem {
    version: 2
    cluster_name: cluster-odoo
    transport: knet
    crypto_cipher: aes256
    crypto_hash: sha256
}

nodelist {
```



```
node {
    ring0_addr: srv1.enermol.lan
    name: srv1.enermol.lan
    nodeid: 1
}

node {
    ring0_addr: srv2.enermol.lan
    name: srv2.enermol.lan
    nodeid: 2
}

quorum {
    provider: corosync_votequorum
    two_node: 1
}

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

```
[root@srv1 ~]# corosync-keygen
Corosync Cluster Engine Authentication key generator.
Gathering 1024 bits for key from /dev/random.
Press keys on your keyboard to generate entropy.
Press keys on your keyboard to generate entropy (bits = 920).
Press keys on your keyboard to generate entropy (bits = 1000).
Writing corosync key to /etc/corosync/authkey.
```

```
[root@srv1 ~]# scp /etc/corosync/authkey srv2:/etc/corosync/
```

```
[root@srv1 ~]# pcs cluster start --all
srv2.enermol.lan: Starting Cluster...
srv1.enermol.lan: Starting Cluster...
```

```
[root@srv1 ~]# pcs cluster enable --all
srv1.enermol.lan: Cluster Enabled
srv2.enermol.lan: Cluster Enabled
```

```
[root@srv1 ~]# pcs status
Cluster name: cluster-odoo
```

WARNINGS:

No stonith devices and stonith-enabled is not false

```
Stack: corosync
Current DC: srv1.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition
with quorum
Last updated: Sun Mar 22 10:06:59 2020
Last change: Sun Mar 22 10:05:41 2020 by hacluster via crmd on srv1.enermol.lan
```

```
2 nodes configured
0 resources configured
```



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

```
No resources
```

```
Daemon Status:
```

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

```
[root@srv1 ~]# corosync-cfgtool -s
```

```
Printing link status.
```

```
Local node ID 1
```

```
LINK ID 0
```

```
addr = 192.168.10.151
```

```
status:
```

```
nodeid 1: link enabled:1 link connected:1  
nodeid 2: link enabled:1 link connected:1
```

```
[root@srv2 ~]# corosync-cfgtool -s
```

```
Printing link status.
```

```
Local node ID 2
```

```
LINK ID 0
```

```
addr = 192.168.10.152
```

```
status:
```

```
nodeid 1: link enabled:1 link connected:1  
nodeid 2: link enabled:1 link connected:1
```

```
[root@srv1 ~]# corosync-cmapctl | grep members
```

```
runtime.members.1.config_version (u64) = 0  
runtime.members.1.ip (str) = r(0) ip(192.168.10.151)  
runtime.members.1.join_count (u32) = 1  
runtime.members.1.status (str) = joined  
runtime.members.2.config_version (u64) = 0  
runtime.members.2.ip (str) = r(0) ip(192.168.10.152)  
runtime.members.2.join_count (u32) = 1  
runtime.members.2.status (str) = joined
```

```
[root@srv1 ~]# journalctl | grep error
```

```
Mar 22 09:03:50 srv1.enermol.lan dracut[6694]: lrwxrwxrwx 1 root root  
22 Jan 3 18:12 usr/lib64/libgpg-error.so.0 -> libgpg-error.so.0.24.2  
Mar 22 09:03:50 srv1.enermol.lan dracut[6694]: -rwxr-xr-x 1 root root  
145984 May 11 2019 usr/lib64/libgpg-error.so.0.24.2  
Mar 22 10:05:41 srv1.enermol.lan pacemaker-schedulerd[5686]: error: Resource  
start-up disabled since no STONITH resources have been defined  
Mar 22 10:05:41 srv1.enermol.lan pacemaker-schedulerd[5686]: error: Either  
configure some or disable STONITH with the stonith-enabled option  
Mar 22 10:05:41 srv1.enermol.lan pacemaker-schedulerd[5686]: error: NOTE:  
Clusters with shared data need STONITH to ensure data integrity  
Mar 22 10:05:41 srv1.enermol.lan pacemaker-schedulerd[5686]: notice:  
Configuration errors found during scheduler processing, please run "crm_verify  
-L" to identify issues
```

```
[root@srv1 ~]# crm_verify -L -V
```

```
(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
```

```
[root@srv1 ~]# pcs status corosync
```

```
Membership information
```

```
-----
Nodeid      Votes Name
   1         1 srv1.enermol.lan (local)
   2         1 srv2.enermol.lan
```

```
[root@srv2 ~]# pcs status corosync
```

```
Membership information
```

```
-----
Nodeid      Votes Name
   1         1 srv1.enermol.lan
   2         1 srv2.enermol.lan (local)
```

```
[root@srv1 ~]# pcs property set stonith-enabled=false
```

```
[root@srv1 ~]# crm_verify -L -V
```

```
[root@srv1 ~]# corosync-quorumtool
```

```
Quorum information
```

```
-----
Date:                Thu Oct 18 18:16:56 2018
Quorum provider:    corosync_votequorum
Nodes:              2
Node ID:            1
Ring ID:            1/20
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 srv1 (local)
   2         1 srv2
```

```
[root@srv1 ~]# pcs quorum status
```

```
Quorum information
```

```
-----
Date:                Sun Mar 22 10:19:21 2020
Quorum provider:    corosync_votequorum
Nodes:              2
Node ID:            1
Ring ID:            1/8
Quorate:            Yes
```

```
Votequorum information
```



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

Membership information

```
-----  
Nodeid      Votes      Qdevice Name  
  1          1          NR srv1.enermol.lan (local)  
  2          1          NR srv2.enermol.lan
```

<https://192.168.10.151:2224>

<https://192.168.10.152:2224>

Username:
Password:



The screenshot shows the 'HIGH AVAILABILITY MANAGEMENT' interface. At the top, there are tabs for 'MANAGE CLUSTERS' and 'PERMISSIONS', and a user dropdown menu showing 'hacluster'. Below the tabs, there is a 'MANAGE CLUSTERS' section with a table showing 0 clusters. To the right, there is an 'INFORMATION ABOUT CLUSTERS' section with the text 'Select a cluster to view more detailed cluster information'. A modal dialog box titled 'Add Existing Cluster' is open, prompting the user to 'Enter the hostname/IP of a node in a cluster that you would like to manage:'. The dialog contains two input fields: 'Node Name/IP:' with the value 'srv1' and 'PCSD port:' with the value '2224'. At the bottom of the dialog are 'Add Existing' and 'Cancel' buttons.

This screenshot shows the same interface as the previous one, but with an additional modal dialog box titled 'Authentication of nodes' open. This dialog prompts the user to 'Enter password for user "hacluster" to authenticate nodes. Nodes to authenticate:'. It shows the node 'srv1' with a password field containing seven dots. At the bottom of this dialog are 'Authenticate' and 'Cancel' buttons. The background interface is dimmed.

MANAGE CLUSTERS Remove Add Existing Create New

NAME	NODES	RESOURCES
cluster-odoo	2	0

INFORMATION ABOUT CLUSTERS

Select a cluster to view more detailed cluster information

Cluster: cluster-odoo hacluster

NODES Remove Add Edit Node

srv1
srv2

Edit Node srv2

srv2 Pacemaker Connected Corosync Connected

Start Stop Restart Standby Maintenance Configure Fencing

Node ID: 2 Uptime: 0 days, 00:47:24

Cluster Daemons

NAME	STATUS
pacemaker	Running (Enabled)
corosync	Running (Enabled)
pcsd	Running (Enabled)

Running Resources

NAME
NONE

Resource Location Preferences

NAME	Score
NONE	

**# netstat -lntp**

Active Internet connections (only servers)

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	0	0.0.0.0:5355	0.0.0.0:*	LISTEN
1382/systemd-resolv					
tcp	0	0	0.0.0.0:111	0.0.0.0:*	LISTEN
1/systemd					
tcp	0	0	0.0.0.0:2224	0.0.0.0:*	LISTEN
4131/platform-pytho					
tcp	0	0	0.0.0.0:22	0.0.0.0:*	LISTEN
969/sshd					
tcp6	0	0	:::5355	:::*	LISTEN
1382/systemd-resolv					
tcp6	0	0	:::111	:::*	LISTEN
1/systemd					
tcp6	0	0	:::2224	:::*	LISTEN
4131/platform-pytho					
tcp6	0	0	:::22	:::*	LISTEN
969/sshd					

ss -nltp

State	Recv-Q	Send-Q	Local Address:Port	Peer Address:Port	
LISTEN	0	128	0.0.0.0:5355	0.0.0.0:*	users:
			(("systemd-resolve", pid=1382, fd=13))		
LISTEN	0	128	0.0.0.0:111	0.0.0.0:*	users:
			(("systemd", pid=1, fd=28))		
LISTEN	0	128	0.0.0.0:2224	0.0.0.0:*	users:
			(("pcsd", pid=4131, fd=6))		
LISTEN	0	128	0.0.0.0:22	0.0.0.0:*	users:
			(("sshd", pid=969, fd=5))		
LISTEN	0	128	:::5355	:::*	users:
			(("systemd-resolve", pid=1382, fd=15))		
LISTEN	0	128	:::111	:::*	users:
			(("systemd", pid=1, fd=30))		
LISTEN	0	128	:::2224	:::*	users:
			(("pcsd", pid=4131, fd=7))		
LISTEN	0	128	:::22	:::*	users:
			(("sshd", pid=969, fd=7))		

firewall-cmd --list-services

cockpit dhcpv6-client high-availability ntp ssh

→ [Instalación de cockpit.](#)**# firewall-cmd --permanent --add-port=9090/tcp**

success

firewall-cmd --reload



```
success
# firewall-cmd --list-ports
123/udp 9090/tcp

# firewall-cmd --list-services
cockpit dhcpv6-client high-availability ntp ssh

# dnf install cockpit* -y

# systemctl enable cockpit.socket --now
Created symlink /etc/systemd/system/sockets.target.wants/cockpit.socket →
/usr/lib/systemd/system/cockpit.socket.

# systemctl status cockpit.socket
● cockpit.socket - Cockpit Web Service Socket
   Loaded: loaded (/usr/lib/systemd/system/cockpit.socket; enabled; vendor
  preset: disabled)
   Active: inactive (dead)
     Docs: man:cockpit-ws(8)
    Listen: [::]:9090 (Stream)

# netstat -nltp
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
PID/Program name
tcp        0      0 0.0.0.0:5355            0.0.0.0:*               LISTEN
1382/systemd-resolv
tcp        0      0 0.0.0.0:111            0.0.0.0:*               LISTEN
1/systemd
tcp        0      0 0.0.0.0:2224           0.0.0.0:*               LISTEN
4131/platform-pytho
tcp        0      0 0.0.0.0:22             0.0.0.0:*               LISTEN
969/sshd
tcp6       0      0 :::5355                :::*                    LISTEN
1382/systemd-resolv
tcp6       0      0 :::111                 :::*                    LISTEN
1/systemd
tcp6       0      0 :::2224                :::*                    LISTEN
4131/platform-pytho
tcp6       0      0 :::22                  :::*                    LISTEN
969/sshd
tcp6       0      0 :::9090                :::*                    LISTEN
1/systemd

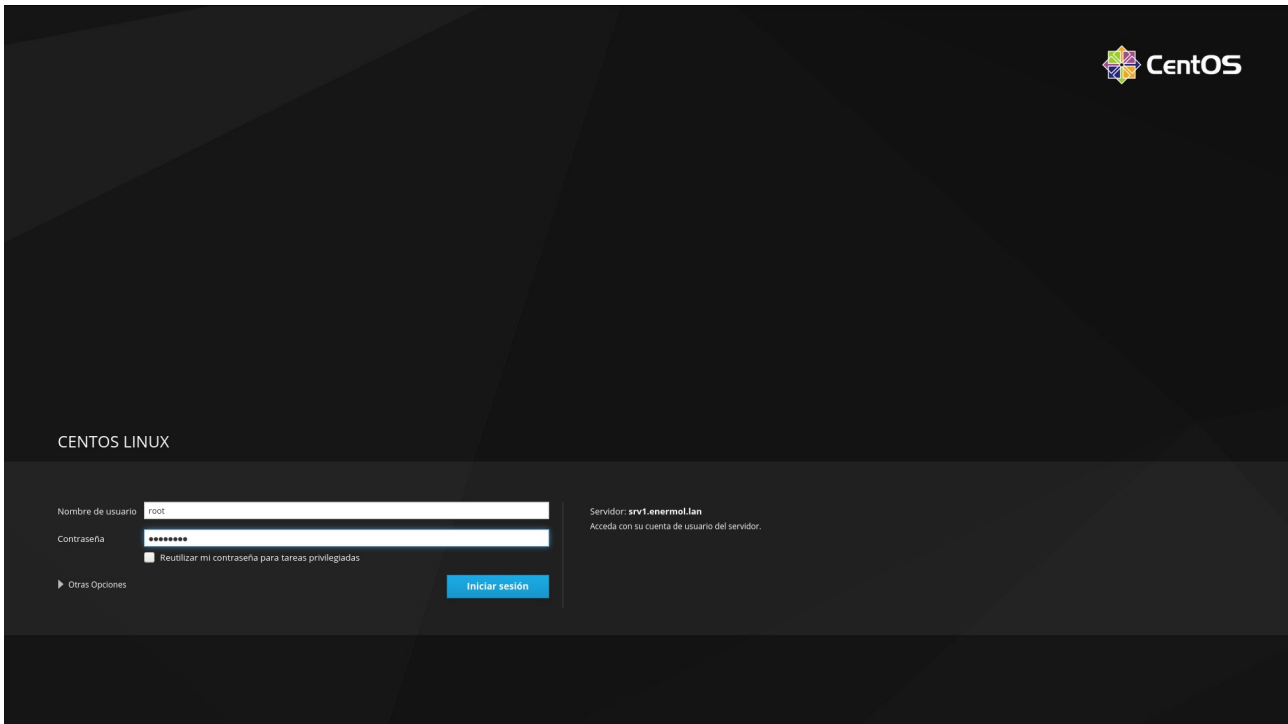
# systemctl status firewalld.service
● firewalld.service - firewalld - dynamic firewall daemon
   Loaded: loaded (/usr/lib/systemd/system/firewalld.service; enabled; vendor
  preset: enabled)
   Active: active (running) since Tue 2020-03-24 11:37:28 CET; 3h 52min ago
     Docs: man:firewalld(1)
  Main PID: 942 (firewalld)
    Tasks: 3 (limit: 17950)
   Memory: 41.4M
   CGroup: /system.slice/firewalld.service
           └─942 /usr/libexec/platform-python -s /usr/sbin/firewalld --nofork --
  nopic

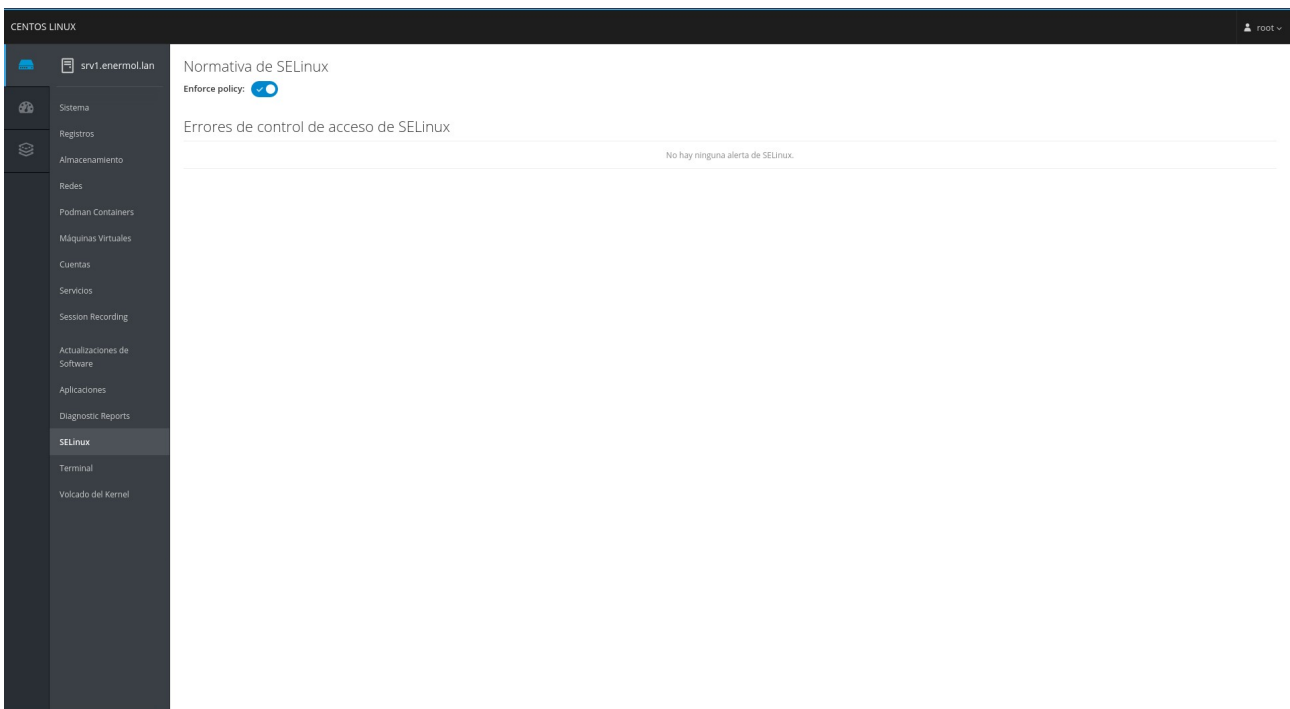
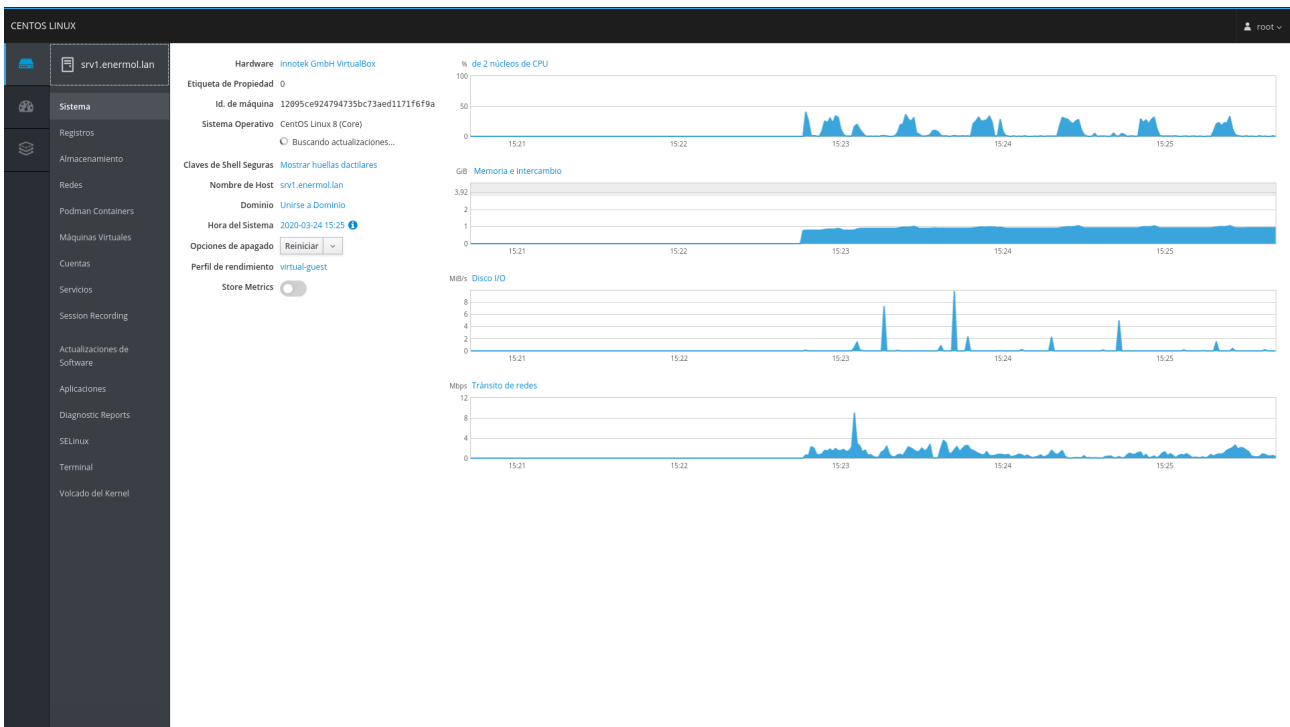
mar 24 11:37:16 srv1.enermol.lan systemd[1]: Starting firewalld - dynamic
```



```
firewall daemon...
mar 24 11:37:28 srv1.enermol.lan systemd[1]: Started firewalld - dynamic
firewall daemon.
mar 24 14:11:08 srv1.enermol.lan firewalld[942]: WARNING: ALREADY_ENABLED:
cockpit
```

<https://192.168.10.151:9090/>
<https://192.168.10.152:9090/>





<https://192.168.99.106:9090/selinux/setroubleshoot>

```
[root@srv1 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group
default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
```



```
inet6 ::1/128 scope host
    valid_lft forever preferred_lft forever
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP
group default qlen 1000
    link/ether 08:00:27:bc:e1:4c brd ff:ff:ff:ff:ff:ff
    inet 192.168.99.106/24 brd 192.168.99.255 scope global dynamic noprefixroute
enp0s3
    valid_lft 346sec preferred_lft 346sec
    inet6 fe80::ba90:c77a:eebc:fb6/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
3: enp0s8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP
group default qlen 1000
    link/ether 08:00:27:f8:1f:e7 brd ff:ff:ff:ff:ff:ff
    inet 192.168.10.151/24 brd 192.168.10.255 scope global noprefixroute enp0s8
    valid_lft forever preferred_lft forever
    inet6 fe80::e7ff:d218:38ad:c237/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
4: enp0s9: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP
group default qlen 1000
    link/ether 08:00:27:2f:45:6f brd ff:ff:ff:ff:ff:ff
    inet 192.168.10.161/24 brd 192.168.10.255 scope global noprefixroute enp0s9
    valid_lft forever preferred_lft forever
    inet6 fe80::b5c0:7a48:7ab4:8c87/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
5: enp0s10: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP
group default qlen 1000
    link/ether 08:00:27:39:01:3a brd ff:ff:ff:ff:ff:ff
    inet 10.0.5.15/24 brd 10.0.5.255 scope global dynamic noprefixroute enp0s10
    valid_lft 69674sec preferred_lft 69674sec
    inet6 fe80::5374:d5df:e304:7a25/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
```

```
[root@srv1 ~]# pcs resource create pgsq1-vip ocf:heartbeat:IPaddr2 nic="enp0s8"
ip=192.168.10.150 cidr_netmask=24 op monitor interval=30s
```

(En realidad el interfaz 'enp0s8' no debería ser el candidato, puesto que ya esta ocupado en otra dirección. Habida cuenta que el modelo presentado se realiza de forma virtualizada y no se dispone de mas interfaces se utiliza el mismo para este propósito, pero teniendo en cuenta que el interfaz en 'modo real' debería ser independiente y único.)

```
[root@srv1 ~]# ip a
```

```
...
3: enp0s8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP
group default qlen 1000
    link/ether 08:00:27:f8:1f:e7 brd ff:ff:ff:ff:ff:ff
    inet 192.168.10.151/24 brd 192.168.10.255 scope global noprefixroute enp0s8
    valid_lft forever preferred_lft forever
    inet 192.168.10.150/24 brd 192.168.10.255 scope global secondary enp0s8
    valid_lft forever preferred_lft forever
    inet6 fe80::e7ff:d218:38ad:c237/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
...
```

```
[root@srv1 ~]# pcs status
Cluster name: cluster-odoo
Stack: corosync
```



```
Current DC: srv2.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition
with quorum
Last updated: Tue Mar 24 16:20:52 2020
Last change: Tue Mar 24 16:18:26 2020 by root via cibadmin on srv1.enermol.lan
```

```
2 nodes configured
1 resource configured
```

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

```
Full list of resources:
```

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

```
Daemon Status:
```

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

```
[root@srv1 ~]# pcs resource create pgsql-alt-vip ocf:heartbeat:IPaddr2
nic="enp0s9" ip=192.168.10.160 cidr_netmask=24 op monitor interval=30s
```

(En realidad y por la misma argumentación anterior el interfaz 'enp0s9' no debería ser el candidato, puesto que ya esta ocupado en otra dirección. Habida cuenta que el modelo presentado se realiza de forma virtualizada y no se dispone de mas interfaces se utiliza el mismo para este propósito, pero teniendo en cuenta que el interfaz en 'modo real' debería ser independiente y único.)

```
# pcs status
```

```
Cluster name: cluster-odoo
Stack: corosync
Current DC: srv1.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition
with quorum
Last updated: Wed Mar 25 09:51:30 2020
Last change: Tue Mar 24 19:03:44 2020 by root via cibadmin on srv1.enermol.lan
```

```
2 nodes configured
2 resources configured
```

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

```
Full list of resources:
```

```
pgsql-vip (ocf::heartbeat:IPaddr2): Started srv1.enermol.lan
pgsql-alt-vip (ocf::heartbeat:IPaddr2): Started srv2.enermol.lan
```

```
Daemon Status:
```

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

```
# ping -c1 pgsql-vip
```

```
PING pgsql-vip.enermol.lan (192.168.10.150) 56(84) bytes of data.
64 bytes from pgsql-vip.enermol.lan (192.168.10.150): icmp_seq=1 ttl=64
time=0.069 ms
```

```
--- pgsql-vip.enermol.lan ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
```



```
rtt min/avg/max/mdev = 0.069/0.069/0.069/0.000 ms
```

```
[root@srv1 ~]# ping -c1 pgsq1-alt-vip
```

```
PING pgsq1-alt-vip.enermol.lan (192.168.10.160) 56(84) bytes of data.  
64 bytes from pgsq1-alt-vip.enermol.lan (192.168.10.160): icmp_seq=1 ttl=64  
time=0.765 ms
```

```
--- pgsq1-alt-vip.enermol.lan ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 0.765/0.765/0.765/0.000 ms
```

```
[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=60
```

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.

```
[root@srv1 ~]# pcs resource defaults
```

```
resource-stickiness=100
```

```
migration-threshold=5
```

```
[root@srv1 ~]# pcs status
```

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

```
Stack: corosync
```

```
Current DC: srv2.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition  
with quorum
```

```
Last updated: Tue Mar 24 17:43:02 2020
```

```
Last change: Tue Mar 24 17:42:28 2020 by root via cibadmin on srv1.enermol.lan
```

```
2 nodes configured
```

```
2 resources configured
```

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

```
Full list of resources:
```

```
pgsq1-vip (ocf::heartbeat:IPaddr2): Started srv1.enermol.lan
```

```
pgsq1-alt-vip (ocf::heartbeat:IPaddr2): Started srv2.enermol.lan
```

```
Daemon Status:
```

```
corosync: active/enabled
```

```
pacemaker: active/enabled
```

```
pcsd: active/enabled
```


[→ Pruebas de HA cruzadas.](#)

```
# pcs node standby srv{1,2}.enermol.lan
# pcs status
Cluster name: cluster-odoo
Stack: corosync
Current DC: srv2.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition
with quorum
Last updated: Tue Mar 24 17:55:48 2020
Last change: Tue Mar 24 17:55:45 2020 by root via cibadmin on srv1.enermol.lan

2 nodes configured
2 resources configured
```

```
Node srv1.enermol.lan: standby
Node srv2.enermol.lan: standby
```

Full list of resources:

```
pgsql-vip (ocf::heartbeat:IPaddr2): Stopped
pgsql-alt-vip (ocf::heartbeat:IPaddr2): Stopped
```

Daemon Status:

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

```
[root@srv1 ~]# pcs resource
```

```
pgsql-vip (ocf::heartbeat:IPaddr2): Stopped
pgsql-alt-vip (ocf::heartbeat:IPaddr2): Stopped
```

```
[root@srv1 ~]# pcs node unstandby srv1.enermol.lan
```

```
[root@srv1 ~]# pcs resource
```

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

```
[root@srv1 ~]# pcs node unstandby srv2.enermol.lan
```

```
[root@srv1 ~]# pcs status
```

```
Cluster name: cluster-odoo
Stack: corosync
Current DC: srv1.enermol.lan (version 2.0.2-3.el8_1.2-744a30d655) - partition
with quorum
Last updated: Wed Mar 25 10:04:36 2020
Last change: Wed Mar 25 10:04:30 2020 by root via cibadmin on srv1.enermol.lan
```

```
2 nodes configured
2 resources configured
```

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

Full list of resources:

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

Daemon Status:

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



```
pcsd: active/enabled
```

```
[root@srv1 ~]# pcs resource move pgsq1-vip srv2.enermol.lan
[root@srv1 ~]# pcs resource move pgsq1-alt-vip srv2.enermol.lan
[root@srv1 ~]# pcs resource
pgsq1-vip (ocf::heartbeat:IPAddr2): Started srv2.enermol.lan
pgsq1-alt-vip (ocf::heartbeat:IPAddr2): Started srv2.enermol.lan
```

```
[root@srv1 ~]# pcs constraint location pgsq1-vip prefers srv1.enermol.lan
[root@srv1 ~]# pcs constraint location
Location Constraints:
Resource: pgsq1-vip
Enabled on:
Node: srv1.enermol.lan (score:INFINITY)
```

→ [Instalar pcp](#) → [Performance Co-Pilot](#).

```
# dnf install pcp-zeroconf -y
# tree /var/log/pcp
/var/log/pcp
├── NOTICES
├── pmcd
│   ├── dm.log
│   ├── kvm.log
│   ├── kvm.log.prev
│   ├── linux.log
│   ├── linux.log.prev
│   ├── nfsclient.log
│   ├── pmcd.log
│   ├── proc.log
│   ├── proc.log.prev
│   ├── root.log
│   ├── root.log.prev
│   ├── xfs.log
│   └── xfs.log.prev
├── pmie
│   └── srv1.enermol.lan
│       └── pmie.log
├── pmlogger
│   ├── pmlogger_check.log
│   ├── pmlogger_check.log.prev
│   ├── pmlogger_daily-K.log
│   └── srv1.enermol.lan
│       ├── 20200805.05.31.0
│       ├── 20200805.05.31.index
│       ├── 20200805.05.31.meta
│       ├── Latest
│       └── pmlogger.log
└── pmproxy
```

```
[root@srv1 ~]# pcs cluster cib configuracion.xml
```

```
[root@srv1 ~]# pcs config backup configuración-cluster0doo.txt
```

```
[root@srv1 ~]# ls configuración-cluster0doo.txt.tar.bz2
```

```
configuración-cluster0doo.txt.tar.bz2
```

**REFERENCIAS:**

<https://www.systutorials.com/docs/linux/man/5-votegorum/>

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

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/configuring_and_managing_high_availability_clusters/index

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.