



Clustering Quick Installation Guide

for PacketFence version 6.7.0

Clustering Quick Installation Guide

by Inverse Inc.

Version 6.5.0 - Jan 2017

Copyright © 2015 Inverse inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

The fonts used in this guide are licensed under the SIL Open Font License, Version 1.1. This license is available with a FAQ at: <http://scripts.sil.org/OFL>

Copyright © Łukasz Dziejdzic, <http://www.latofonts.com>, with Reserved Font Name: "Lato".

Copyright © Raph Levien, <http://levien.com/>, with Reserved Font Name: "Inconsolata".

9279VnJ

Table of Contents

| | |
|---|----|
| About this Guide | 1 |
| Assumptions | 2 |
| Installation on CentOS 6 | 3 |
| Step 1: Install the replicated database | 3 |
| Step 2 : Server configuration | 9 |
| Installation on CentOS 7 | 11 |
| Step 1: Install the replicated database | 11 |
| Step 2 : Server configuration | 17 |
| Step 3 : Create a new cluster | 20 |
| Step 4 : Connect a slave packetfence server | 22 |
| Advanced configuration | 24 |
| Removing a server from the cluster | 24 |
| Resynchronizing the configuration manually | 24 |
| Adding files to the synchronization | 25 |
| haproxy dashboard | 25 |
| Configuration conflict handling | 26 |
| Unsupported features in active/active | 27 |
| Appendix | 28 |
| Setting the interfaces name on CentOS 7 | 28 |

About this Guide

This guide has been created to give a quick start to install active/active clustering in PacketFence 5+. This guide does not include advanced troubleshooting of the active/active clustering. Refer to the documentation of HAProxy and Keepalive for advanced features.

Assumptions

- You have at least two (2) installed and configured PacketFence (v5+) servers
- Servers are running RHEL / CentOS 6 or 7
- Servers have identical copies for network interfaces
- Servers network interfaces are on the same layer 2 network
- Servers hostnames must be resolvable via DNS resolution
- Servers have an empty, unformatted partition (same size) that will be used for the database replication

Installation on CentOS 6

Step 1: Install the replicated database



Note

In this example, the database is replicated in active/passive across two servers that are part of the PacketFence cluster. Active/active replication is also possible using MariaDB Galera cluster but this subject is not covered in this guide.

Creation of the DRBD partition

During the OS installation, reduce the size of the main partition by 30G. **Do not** create that partition during the install process, we will do it later.

Partitioning

After the installation, you need to create the extra partition for DRBD. Using fdisk, create your new partition and save the table. You may have to reboot your server after this step if requested.

DRBD and Linux-HA Installation

Add the repository ELRepo.

```
# yum localinstall http://www.elrepo.org/elrepo-release-6-6.el6.elrepo.noarch.rpm
```

Edit the repo file to disable ELRepo by default:

```
/etc/yum.repos.d/elrepo.repo
```

```
[elrepo]
name=ELRepo.org Community Enterprise Linux Repository - el6
baseurl=http://elrepo.org/linux/elrepo/el6/$basearch/
mirrorlist=http://elrepo.org/mirrors-elrepo.el6
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-elrepo.org
protect=0
```

Install DRBD v8.4:

```
# yum install kmod-drbd84 --enablerepo=elrepo
```

DRBD Configuration and setup



Caution

Initializing, configuring and troubleshooting DRBD is not straight forward! We strongly recommend that you read the online documentation available from the DRBD website so you have a better idea about how it works.

Here we assume the name of the partition is mysql.

Load the DRBD kernel module:

```
# modprobe drbd
```

Edit `/etc/drbd.d/global_common.conf` with the following content:

```
global {
    usage-count yes;
}

common {
    protocol C;

    startup {
        degr-wfc-timeout 120;
    }

    syncer {
        rate 100M;
        al-extents 257;
    }

    disk {
        on-io-error    detach;
    }
}
```

Create the file `/etc/drbd.d/mysql.res` with the following content:

```
resource mysql {
  on <pf1_server_name> {
    device /dev/drbd0;
    disk <storage_device>;
    meta-disk internal;
    address <ha_interface_ip_address_1>:7788;
  }

  on <pf2_server_name> {
    device /dev/drbd0;
    disk <storage_device>;
    meta-disk internal;
    address <ha_interface_ip_address_2>:7788;
  }
  handlers {
    split-brain "/usr/lib/drbd/notify-split-brain.sh alert@acme.com";
  }
}
```

where:

- `mysql` is the name of the partition you created when installing the OS
- `pf1_server_name` and `pf2_server_name` by the real server names ([FQDN](#))
- `ha_interface_ip_address_1` and `ha_interface_ip_address_2` by the IP addresses of the management interface on both servers.
- `storage_device` is the device to use for the MySQL partition (ie. `/dev/sda2`)

Then initialize the partition:

```
[root@pf1 ~]# drbdadm create-md mysql
Writing meta data...
initializing activity log
NOT initialized bitmap
New drbd meta data block successfully created.
success
```

Start DRBD on both servers:

```
# /etc/init.d/drbd start
```

You should see something similar to this when typing `cat /proc/drbd`:

```
...
0: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent C r----
ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:30702640
```



Note

If you get connectivity issues make sure your iptables are disabled.

Synchronize the servers by forcing one to become the primary. So on pf1 do:

```
# drbdadm primary --force mysql
```

After issuing this command, the initial full synchronization will start. You will be able to monitor its progress via `/proc/drbd`. It may take some time depending on the size of the device. Wait until it completes.

When the sync is complete, create the filesystem on the primary node only:

```
# mkfs.ext4 /dev/drbd0
```

Make sure DRBD is started at boot time:

```
# chkconfig --level 2345 drbd on
```

Restart both servers.

When done, look in `/proc/drbd` and make sure you see:

```
...
0: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r---
   ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:0
```

MySQL Configuration

First you need to ensure it will not start on boot :

```
``` # chkconfig mysql off ```
```



### Note

By default MySQL puts its data in `/var/lib/mysql`. In order to replicate data between the two servers, we mount the DRBD partition under `/var/lib/mysql`.

When first starting MySQL, the partition must be mounted.

In order to do so:

On the master server (the server you are working on), tell DRBD to become the primary node with:

```
drbdadm primary mysql
```

`mysql` being the name of the DRBD partition.

The command `cat /proc/drbd` should display something like:

```
...
0: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r----
 ns:145068 nr:4448 dw:149516 dr:10490 al:31 bm:14 lo:0 pe:0 ua:0 ap:0 ep:1
 wo:d oos:0
```

Stop and backup previous MySQL data

```
service mysqld stop
mkdir /var/lib/mysql.bak
mv /var/lib/mysql/* /var/lib/mysql.bak/ #Optional
```

Mount the partition with:

```
mount /dev/drbd0 /var/lib/mysql
```

Restore previous MariaDB data

```
mv /var/lib/mysql.bak/* /var/lib/mysql/ #Optional
```

Start MySQL

```
service mysqld start
```

Execute the secure installation script in order to set the root password, remove the test databases and anonymous user created by default:

```
/usr/bin/mysql_secure_installation
```

Make sure MySQL does **not** start at boot time:

```
chkconfig --level 2345 mysqld off
```

## Heartbeat configuration

Install the EPEL repository

```
yum localinstall http://dl.fedoraproject.org/pub/epel/6/x86_64/epel-release-6-8.noarch.rpm
```

In `/etc/yum.repos.d/epel.repo`, disable the repository

```
[epel]
name=Extra Packages for Enterprise Linux 6 - $basearch
#baseurl=http://download.fedoraproject.org/pub/epel/6/$basearch
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=epel-6&arch=$basearch
failovermethod=priority
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-6
```

Now install Heartbeat

```
yum install heartbeat --enablerepo=epel
```

Create `/etc/ha.d/ha.cf` with the following content:

```
bcast eth0
keepalive 2
warntime 30
deadtime 60
auto_failback off
initdead 120
node pf1.example.com
node pf2.example.com
use_logd yes
logfile /var/log/ha-log
```

Here we assume that the connection for the Heartbeat between the 2 servers is on `eth0`. Make sure you change `pf1.example.com` and `pf2.example.com` to the hostnames of your servers.

Create `/etc/ha.d/haresources` with the following content:

```
pf1.example.com drbdisk::mysql Filesystem::/dev/drbd0::/var/lib/mysql::ext4 \
mysqld
```

As for the previous file, make sure to adjust hostname, in this case `pf1.example.com`, accordingly for each servers.

Create `/etc/ha.d/authkeys` with the following content (make sure to change the hash for security concerns):

```
auth 1
1 sha1 10abf064f24a3e807dde7b945d0303392f10777d
```

and change its rights like this:

```
chmod 600 /etc/ha.d/authkeys
```

Changing the hash can be achieved by running `echo -n password | sha1sum | awk '{print $1}'`, where `password` would be a password of your choice.



### Note

Make sure udp port 694 is opened (through iptables) on both servers. Important mainly if an existing iptables configuration is in place on the server. PacketFence does allow it on it's *high-availability* interface.

Start Heartbeat:

```
service heartbeat start
```

Look at Heartbeat log file `/var/log/ha-log` to make sure that everything is fine.

Enable HB automatic start

```
chkconfig --level 345 heartbeat on
```

## MySQL network configuration

In order for your PacketFence cluster to communicate properly with the MySQL database, you need to have MySQL bind on the management IP address.

Adjust your MySQL configuration in `/etc/my.cnf` on both servers and make sure the `bind_address` line is :

```
bind_address=<management_ip_address>
```

## Step 2 : Server configuration

---

You will need to configure the server so the services can bind on IP addresses they don't currently have configured. This allows faster failover of the services.

On CentOS, add the following line in `/etc/sysctl.conf` and then reload with `sysctl -p`

```
net.ipv4.ip_nonlocal_bind = 1
```

Create the PEM that combines the key and certificate for the http services. Adapt to your own paths if you are using different certificates.

```
cd /usr/local/pf/conf/ssl
cat server.key server.crt > server.pem
```

## MySQL configuration

In order for PacketFence to communicate properly with your MySQL cluster, you need to change the following.

In `conf/pf.conf` :

```
[database]
host=127.0.0.1

[monitoring]
db_host=127.0.0.1
....
```

In `conf/pfconfig.conf` :

```
[mysql]
host=127.0.0.1
....
```

In `/usr/local/fingerbank/conf/fingerbank.conf`

```
[mysql]
host=127.0.0.1
```

Make sure you restart MySQL, packetfence-config and packetfence

```
service mysqld restart
service packetfence-config restart
service packetfence restart
```

Next, you need to remove the empty users from your MySQL database

```
mysql
mysql> delete from mysql.user where user = '' ;
mysql> flush privileges;
```

# Installation on CentOS 7

---

## Step 1: Install the replicated database

---



### Note

In this example, the database is replicated in active/passive across two servers that are part of the PacketFence cluster. Active/active replication is also possible using MariaDB Galera cluster but this subject is not covered in this guide.

## Creation of the DRBD partition

During the OS installation, reduce the size of the main partition by 30G. **Do not** create that partition during the install process, we will do it later. This needs to be done on both servers.

## Partitioning

After installation, you need to create the extra partition for DRBD. Using `fdisk`, create your new partition and save the table. You will probably need to reboot your server after this step. This also needs to be done on both servers.

## Clustering prerequisites

Make sure that you systems are correctly configured and up-to-date

```
yum update
```

`/etc/hosts` must contain the following two lines:

```
ha_interface_ip_address_1 pf1_server_name
ha_interface_ip_address_2 pf2_server_name
```

Set hostname on each server:



## Note

An active NetworkManager will overwrite your DNS configuration if not explicitly instructed to not do so before executing the following commands. ([Bug Report](#)) ([Corrective Measures](#))

```
hostnamectl set-hostname pf1_server_name
hostnamectl set-hostname pf2_server_name
```

Disable selinux on each server:

```
setenforce 0
sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
```

- pf1\_server\_name and pf2\_server\_name by the real server names ([FQDN](#))
- ha\_interface\_ip\_address\_1 and ha\_interface\_ip\_address\_2 by the IP addresses of the management interface on both servers.

## Pacemaker and Corosync installation and configuration

Install pcs (pacemaker/corosync) configuration system

```
yum install pcs
```

First you need to disable firewalld and replace it by iptables

```
systemctl disable firewalld
yum install iptables-services
systemctl enable iptables.service
iptables -F
iptables -X
iptables -t nat -F
iptables -t nat -X
iptables -t mangle -F
iptables -t mangle -X
iptables -P INPUT ACCEPT
iptables -P FORWARD ACCEPT
iptables -P OUTPUT ACCEPT
service iptables save
```

Next you need to set up a password for the pcs administration account named hacluster

```
echo "passwd" | passwd hacluster --stdin
```

Start and enable the service

```
systemctl start pcsd.service
systemctl enable pcsd.service
```

Enable corosync and pacemaker

```
systemctl enable corosync
systemctl enable pacemaker
```

Create a new cluster:

```
pcs cluster auth pf1_server_name pf2_server_name -u hacluster -p passwd
pcs cluster setup --name mysql_cluster pf1_server_name pf2_server_name --token
 15000 --join 1000 --consensus 6000 -u hacluster -p passwd
pcs cluster start --all
```

## DRBD Configuration and setup

Add the repository ELRepo.

```
yum localinstall http://www.elrepo.org/elrepo-
 release-7.0-2.el7.elrepo.noarch.rpm
```

Edit the repo file to disable ELRepo by default:

```
/etc/yum.repos.d/elrepo.repo
```

```
[elrepo]
name=ELRepo.org Community Enterprise Linux Repository - el7
baseurl=http://elrepo.org/linux/elrepo/el7/$basearch/
 http://mirrors.coreix.net/elrepo/elrepo/el7/$basearch/
 http://jur-linux.org/download/elrepo/elrepo/el7/$basearch/
 http://repos.lax-noc.com/elrepo/elrepo/el7/$basearch/
 http://mirror.ventraip.net.au/elrepo/elrepo/el7/$basearch/
mirrorlist=http://mirrors.elrepo.org/mirrors-elrepo.el7
enabled=0
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-elrepo.org
protect=0
```

Install now the package DRBD v8.4 enabling .

```
yum install kmod-drbd84 --enablerepo=elrepo
```



### Caution

Initializing, configuring and troubleshooting DRBD is not straight forward! We strongly recommend that you read the online documentation available from the DRBD website so you have a better idea about how it works.

Here we assume the name of the partition is mysql.



Load the DRBD kernel module:

```
modprobe drbd
```

Create the file `/etc/drbd.d/global_common.conf`:

```
cat << EOL >/etc/drbd.d/global_common.conf
global {
 usage-count yes;
}

common {
 protocol C;

 startup {
 degr-wfc-timeout 120;
 }

 syncer {
 rate 100M;
 al-extents 257;
 }

 disk {
 on-io-error detach;
 }
}
EOL
```

Create the file `/etc/drbd.d/mysql.res`:

```

cat << EOL >/etc/drbd.d/mysql.res
resource mysql {
 protocol C;
 meta-disk internal;
 device /dev/drbd0;
 disk storage_device;
 handlers {
 split-brain "/usr/lib/drbd/notify-split-brain.sh root";
 }
 net {
 allow-two-primaries no;
 after-sb-0pri discard-zero-changes;
 after-sb-2pri disconnect;
 rr-conflict disconnect;
 }
 disk {
 on-io-error detach;
 }
 syncer {
 verify-alg sha1;
 }
 on pf1 {
 address ha_interface_ip_address_1:7788;
 }
 on pf2 {
 address ha_interface_ip_address_2:7788;
 }
}
EOL

```

where:

- `ha_interface_ip_address_1` and `ha_interface_ip_address_2` by the IP addresses of the management interface on both servers.
- `pf1` and `pf2` are the hostname of your servers.
- `storage_device` is the device to use for the MySQL partition (ie. `/dev/sda2`)

Then initialize the partition on both servers:

```

[root@pf1 ~]# drbdadm create-md mysql
Writing meta data...
initializing activity log
NOT initialized bitmap
New drbd meta data block successfully created.
success

```

```

[root@pf2 ~]# drbdadm create-md mysql
Writing meta data...
initializing activity log
NOT initialized bitmap
New drbd meta data block successfully created.
success

```

Start DRBD on both servers:

```
systemctl start drbd
```

You should see something similar to this when typing `cat /proc/drbd`:

```
...
0: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent C r----
 ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:30702640
```



### Note

If you get connectivity issues make sure your iptables are disabled.

Synchronize the servers by forcing one to become the primary. So on pf1 do:

```
drbdadm primary --force mysql
```

After issuing this command, the initial full synchronization will start. You will be able to monitor its progress via `/proc/drbd`. It may take some time depending on the size of the device. Wait until it completes.

When the sync is complete, create the filesystem on the primary node only:

```
mkfs.ext4 /dev/drbd0
```

## MariaDB Configuration

First you need to ensure it will not start on boot :

```
`` # systemctl disable mariadb ``
```



### Note

By default MariaDB puts its data in `/var/lib/mysql`. In order to replicate data between the two servers, we mount the DRBD partition under `/var/lib/mysql`.

When first starting MariaDB, the partition must be mounted.

In order to do so:

On the master server (the server you are working on), tell DRBD to become the primary node with:

```
drbdadm primary mysql
```

`mysql` being the name of the DRBD partition.

The command `cat /proc/drbd` should display something like:

```
...
0: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r----
 ns:145068 nr:4448 dw:149516 dr:10490 al:31 bm:14 lo:0 pe:0 ua:0 ap:0 ep:1
 wo:d oos:0
```

Stop and backup previous MariaDB data

```
systemctl stop mariadb
mkdir /var/lib/mysql.bak
mv /var/lib/mysql/* /var/lib/mysql.bak/ #Optional
```

Mount the partition with:

```
mount /dev/drbd0 /var/lib/mysql
```

Restore previous MariaDB data

```
mv /var/lib/mysql.bak/* /var/lib/mysql/ #Optional
```

Start MariaDB

```
systemctl start mariadb
```

Execute the secure installation script in order to set the root password, remove the test databases and anonymous user created by default:

```
/usr/bin/mysql_secure_installation
```

Make sure to stop MariaDB:

```
systemctl stop mariadb
```

## MariaDB network configuration

In order for your PacketFence cluster to communicate properly with the MariaDB database, you need to have MariaDB bind on the management IP address.

Adjust your MariaDB configuration in `/etc/my.cnf` on both servers and make sure the `bind_address` line is :

```
bind_address=<management_ip_address>
```

## Step 2 : Server configuration

---

You need to make sure the interfaces name will be the same on both servers. See *Setting the interfaces name on CentOS 7* in the Appendix section of this document.

You will need to configure the server so the services can bind on IP addresses they don't currently have configured. This allows faster failover of the services.

On CentOS, add the following line in `/etc/sysctl.conf` and then reload with `sysctl -p`

```
net.ipv4.ip_nonlocal_bind = 1
```

Create the PEM that combines the key and certificate for the http services. Adapt to your own paths if you are using different certificates.

```
cd /usr/local/pf/conf/ssl
cat server.key server.crt > server.pem
```

## Corosync configuration

```
pcs cluster start --all
pcs cluster cib clust_cfg
pcs -f clust_cfg property set stonith-enabled=false
pcs -f clust_cfg property set no-quorum-policy=ignore
pcs -f clust_cfg resource defaults resource-stickiness=200
pcs -f clust_cfg resource create DRBD ocf:linbit:drbd \
 drbd_resource=mysql \
 op monitor interval=30s
pcs -f clust_cfg resource master MySQLClone DRBD \
 master-max=1 master-node-max=1 \
 clone-max=2 clone-node-max=1 \
 notify=true
pcs -f clust_cfg resource create DRBD_fs Filesystem \
 device="/dev/drbd0" \
 directory="/var/lib/mysql" \
 fstype="ext4"
```

```
pcs -f clust_cfg constraint colocation add DRBD_fs with MySQLClone \
 INFINITY with-rsc-role=Master
```

```
pcs -f clust_cfg constraint order promote MySQLClone then start DRBD_fs
```

```
pcs -f clust_cfg resource create MariaDB systemd:mariadb \
 op start timeout=60s \
 op stop timeout=60s \
 op monitor interval=20s timeout=30s
```

```
pcs -f clust_cfg constraint colocation add MariaDB with DRBD_fs INFINITY
```

```
pcs -f clust_cfg constraint order DRBD_fs then MariaDB
```

```
pcs cluster cib-push clust_cfg
```

Check the status of the cluster:

```
pcs status
```

Then try a failover:

```
pcs cluster standby pf1_server_name
```

```
pcs status
```

```
``` Cluster name: mysql_cluster Last updated: Tue Apr 19 09:38:56 2016 Last change: Tue Apr
19 09:38:47 2016 by root via crm_attribute on pf1 Stack: corosync Current DC: pf2 (version
1.1.13-10.el7_2.2-44eb2dd) - partition with quorum 2 nodes and 4 resources configured
```

```
Node pf1_server_name: standby Online: [ pf2_server_name ]
```

Full list of resources:

```
Master/Slave Set: MySQLClone [DRBD]
  Masters: [ pf2 ]
  Stopped: [ pf1 ]
DRBD_fs      (ocf::heartbeat:Filesystem):    Started pf2_server_name
MariaDB      (systemd:mariadb):          Started pf2_server_name
```

```
PCSD Status: pf1: Online pf2: Online
```

```
Daemon Status: corosync: active/disabled pacemaker: active/disabled pcsd: active/enabled ```
```

```
# pcs cluster unstandby pf1_server_name
```

On reboot of one of the members

In the case of a reboot of one of the server you need to launch:

```
# pcs cluster start
```

PacketFence Configuration Modification

In order for PacketFence to communicate properly with your MySQL cluster, you need to change the following.

In conf/pf.conf :

```
[database]
host=127.0.0.1

[monitoring]
db_host=127.0.0.1
....
```

While you are in that file, add "high-availability" statement at the end of your management interface, to be consistent with what will be found in cluster.conf:

```
[interface eth0]
....
type=management,high-availability
....
```

In conf/pfconfig.conf :

```
[mysql]
host=127.0.0.1
....
```

Make sure you restart MySQL:

```
# systemctl restart mariadb
```

Next, you need to remove the empty users from your MySQL database

```
# mysql
mysql> delete from mysql.user where user = '' ;
mysql> flush privileges;
```

Step 3 : Create a new cluster



Note

It is highly recommended to modify the keepalive shared secret in your cluster to prevent attacks. In the administration interface, go in *Configuration/Clustering* and change the *Shared KEY*. Make sure you restart keepalived on both your servers using `/usr/local/pf/bin/pfcmd service keepalived restart`

In order to create a new cluster, the only thing needed is to configure `/usr/local/pf/conf/cluster.conf`

You will need to configure it with your server hostname. To get it use : **hostname** in a command line.

In the case of this example it will be `pf1.example.com`.

The *CLUSTER* section represents the virtual IP addresses of your cluster that will be shared by your servers.

In this example, eth0 is the management interface, eth1.2 is the registration interface and eth1.3 is the isolation interface.

Create a configuration similar to this :

```
[CLUSTER]
management_ip=192.168.1.10

[CLUSTER interface eth0]
ip=192.168.1.10
type=management,high-availability

[CLUSTER interface eth1.2]
ip=192.168.2.10
type=internal

[CLUSTER interface eth1.3]
ip=192.168.3.10
type=internal
```

```
[pf1.example.com]
management_ip=192.168.1.5

[pf1.example.com interface eth0]
ip=192.168.1.5
type=management,high-availability
mask=255.255.255.0

[pf1.example.com interface eth1.2]
enforcement=vlan
ip=192.168.2.5
type=internal
mask=255.255.255.0

[pf1.example.com interface eth1.3]
enforcement=vlan
ip=192.168.3.5
type=internal
mask=255.255.255.0
```



```
[pf2.example.com]
management_ip=192.168.1.6

[pf2.example.com interface eth0]
ip=192.168.1.6
type=management,high-availability
mask=255.255.255.0

[pf2.example.com interface eth1.2]
enforcement=vlan
ip=192.168.2.6
type=internal
mask=255.255.255.0

[pf2.example.com interface eth1.3]
enforcement=vlan
ip=192.168.3.6
type=internal
mask=255.255.255.0
```

Once this configuration is done, restart PacketFence to have it applied:

```
# systemctl restart packetfence
```

If no error is found in the previous configuration, the previous restart of packetfence should have started two additional services : haproxy and keepalived

You should now have service on the IP addresses defined in the *CLUSTER* sections

Step 4 : Connect a slave packetfence server

First, connect the server to the database cluster using the instructions above

On CentOS, add the following line in `/etc/sysctl.conf` and then reload with `sysctl -p`

```
net.ipv4.ip_nonlocal_bind = 1
```

Go through the configurator to setup the interfaces and then stop

Get the webservices user and password on the master node in *Configuration/Web Services* If there's none, set the user, password and then restart `httpd.webservices`

Do (and make sure it does it without any errors) :

```
# /usr/local/pf/bin/cluster/sync --from=192.168.1.5 --api-user=packet --api-
password=fence
```

Where :

- `192.168.1.5` is the management IP of the other PacketFence node
- `packet` is the webservices username on the master node
- `fence` is the webservices password on the master node

Reload the configuration and start the webservices on the new server

```
# service packetfence-config restart
# /usr/local/pf/bin/pfcmd configreload
# /usr/local/pf/bin/pfcmd service haproxy restart
# /usr/local/pf/bin/pfcmd service httpd.webservices restart
```

Make sure that this server is binding to it's own management address. If it's not, verify the `/usr/local/pf/conf/cluster.conf` management interface configuration.

```
# netstat -nlp | grep 9090
```

Now replicate this server configuration to the other nodes in the cluster

```
# /usr/local/pf/bin/cluster/sync --as-master
```

Make sure at least `/usr/local/pf/conf/cluster.conf` was replicated to the other servers

Make sure to join domains through *Configuration/Domains* on each additional create node.

Now restart packetfence on each cluster server keeping the new node as the last one to be restarted.

Advanced configuration

Removing a server from the cluster



Note

Removing a server from the cluster requires a restart of the PacketFence service on all nodes.

First, you will need to stop PacketFence on your server and put it offline.

```
# service packetfence stop
# shutdown -h 0
```

Then you need to remove all the configuration associated to the server from `/usr/local/pf/conf/cluster.conf` on one of the remaining nodes. Configuration for a server is always prefixed by the server's hostname.

Once you have removed the configuration, you need to reload it and synchronize it with the remaining nodes in the cluster.

```
# /usr/local/pf/bin/pfcmd configreload hard
# /usr/local/pf/bin/cluster/sync --as-master
```

Now restart PacketFence on all the servers so that the removed node is not part of the clustering configuration.

Resynchronizing the configuration manually

If you did a manual change in a configuration file, an additional step is now needed.

In order to be sure the configuration is properly synced on all nodes, you will need to enter this command on the previously selected master node.

```
# /usr/local/pf/bin/cluster/sync --as-master
```

Adding files to the synchronization

In the event that you do modifications to non-synchronized files like switch modules, files in raddb/, etc, you can add those files to be synchronized when using `/usr/local/pf/bin/cluster/sync`.

On one of the nodes, create `/usr/local/pf/conf/cluster-files.txt`

Add the additional files one per line in this file. We advise you add this file to the synchronization too.

Example :

```
/usr/local/pf/conf/cluster-files.txt
/usr/local/pf/raddb/modules/mschap
```

haproxy dashboard

You have the possibility to configure the haproxy dashboard which will give you statistics about the current state of your cluster.

In order to active it uncomment the following lines from `/usr/local/pf/conf/haproxy.conf`

```
listen stats %%active_active_ip%:1025
  mode http
  timeout connect 10s
  timeout client 1m
  timeout server 1m
  stats enable
  stats uri /stats
  stats realm HAProxy\ Statistics
  stats auth admin:packetfence
```



Note

We strongly advise you change the username and password to something else than `admin/packetfence` although access to this dashboard doesn't compromise the server.

Next, uncomment the following line from `/usr/local/pf/conf/iptables.conf`



Caution

If you're upgrading from a version prior to 5.0, the line may not be there. Add it close to the other management rules

```
-A input-management-if --protocol tcp --match tcp --dport 1025 --jump ACCEPT
```

Now restart haproxy and iptables in order to complete the configuration

```
# /usr/local/pf/bin/pfcmd service haproxy restart
# /usr/local/pf/bin/pfcmd service iptables restart
```

You should now be able to connect to the dashboard on the following URL : <http://pf.local:1025/stats>

Configuration conflict handling



Note

It is not recommended to perform configuration while one or more node of the cluster is experiencing issues. Still, should that be the case, this section will explain the conflict resolution that will occur when the nodes will reattach together.

When modifying the configuration through the administration interface, the configuration will be automatically synchronized to all the nodes that are online. In the event that one or more nodes cannot be updated, an error message will be displayed with affected nodes.

A scheduled check runs on the management server (controlled through `maintenance.cluster_check_interval`) in order to validate if all servers are running the same configuration version. When the failed node(s) will come back online, that scheduled check will ensure that the new configuration is pushed on the new node(s). You can disable this check by setting `maintenance.cluster_check_interval` to 0 and restarting pfmon. In that case, you will need to manually resolve the conflict when the node(s) come back online by running `/usr/local/pf/bin/cluster/sync --as-master` on the node you want to keep the configuration of.

General facts about conflict resolution:

- If the configuration is not pushed to at least half of the servers of your cluster, when the failed nodes will come back online, they will have quorum on the previous configuration and the one they are running will be pushed to all the servers.
- In a two node cluster, the most recent configuration is always selected when resolving a conflict.
- In a two node cluster, no decision is taken unless the peer server has its webservice available.

Going deeper in the conflict handling

The section below will explain with more details, the steps that are taken in order to take the decision of which server should be declared as the master when one or more servers have conflicting configuration version.

The first step is to get the configuration version from each server through a webservice call.

The results are then organized by version identifier. Should all alive servers run the same version, the state is considered as healthy and nothing happens.

Then, should there be more than one version identifier across the alive servers, the algorithm validates that there are at least 2 servers **configured** in the cluster. If there aren't, then the most recent version is pushed on the peer node.

After that, the algorithm looks at which version is on the most servers. In the event that the dead servers are in higher number than the alive ones, the most recent version is taken. Otherwise, the version that is present on the most servers will be selected.

When pushing a version to the other servers, if the current version has the most recent version or is part of the majority (depending on which push strategy was defined above), then he will be the one pushing the new configuration to the other servers. Otherwise, a webservice call is made to one of the servers running the selected version so that he pushes its configuration to its peers.

Unsupported features in active/active

The following features are not supported when using active/active clustering.

- Switches using SNMP based enforcement (port-security, link up/down, ...)
- SNMP roaming with the Aerohive controller. (4.7+ includes support for accounting roaming)
- Inline enforcement

Appendix

Setting the interfaces name on CentOS 7

On CentOS 7 you need to make sure that all the servers in the cluster use the same interfaces name. This section covers how to set the interfaces to the ethX format. Note that you can set it to the format you desire as long as the names are the same on all servers.

First, go in `/etc/default/grub` and add `net.ifnames=0` to the variable: `GRUB_CMDLINE_LINUX`.

```
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=centos/root rd.lvm.lv=centos/swap
rhgb quiet net.ifnames=0"
```

Then regenerate the GRUB configuration by executing the following command:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Then, rename the network script of your management interface (`eno16780032` in this example) to be in the ethX form (`eth0` in this example)

```
# mv /etc/sysconfig/network-scripts/ifcfg-eno16780032 /etc/sysconfig/network-
scripts/ifcfg-eth0
```

And rename the name of the interface in the following files (making sure you replace `eno16780032` and `eth0` by the appropriate values):

```
# sed -i.bak "s/eno16780032/eth0/g" /etc/sysconfig/network-scripts/ifcfg-eth0
# sed -i.bak "s/eno16780032/eth0/g" /usr/local/pf/conf/pf.conf
```

Apply the last two steps for any other interface you have on your server. Keep in mind, you can use the PacketFence configurator to reconfigure them later as long as your management interface is correctly configured and is accessible on your network.

Now, reboot your server and when it finishes starting, your interfaces name should now be using the format `ethX`