



PacketFence Inline Deployment Quick Guide using ZEN

for PacketFence version 5.3.0

PacketFence Inline Deployment Quick Guide using ZEN

by Inverse Inc.

Version 5.3.0 - July 2015

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 Dziejczak, <http://www.latofonts.com>, with Reserved Font Name: "Lato".

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



Table of Contents

About this Guide	1
Other sources of information	1
Getting Started	2
Virtual Machine	2
Inline	2
Assumptions	3
Network Setup	3
DHCP/DNS	3
Installation	4
Import the virtual machine	4
Virtual Machine passwords	5
Configuration	6
Inline enforcement configuration	6
PacketFence configuration files	9
Network Devices	9
FreeRADIUS	10
Inline Access	10
OMAPI	10
Traffic shaping	11
Test	14
Register a device in inline enforcement	14
Additional Information	15
Commercial Support and Contact Information	16
GNU Free Documentation License	17

About this Guide

This guide will walk you through the installation and configuration of the PacketFence ZEN solution. It covers Inline isolation setup.

The instructions are based on version 5.3.0 of PacketFence.

The latest version of this guide is available online at <http://www.packetfence.org/documentation/guides.html>

Other sources of information

We suggest that you also have a look in the PacketFence Administration Guide, and in the PacketFence Network Devices Configuration Guide. Both are available online at <http://www.packetfence.org/documentation/guides.html>

Getting Started

Virtual Machine

This setup has been tested using VMWare ESXi 4.0 & 5.0 with 8GB of RAM dedicated to the virtual machine. It might work using other virtualization products. You need to have a 64-bit capable CPU on your host.

Inline

In order to build an Inline setup you need :

- 2 network interfaces for the VM (1 for the Inline and another one to go out).
- a switch port in the management network for the PacketFence ZEN box (for eth0).
- a switch port in the inline network for the PacketFence ZEN box (for eth1) which needs to be configured in access mode and in the same access VLAN as every switchport on which devices will be connected.
- your server has net.ipv4.ip_forward enable. Edit the following file:

```
# /etc/sysctl.conf
```

Change net.ipv4.ip_forward from 0 to 1

```
# net.ipv4.ip_forward = 1
```

Now you need to make this change permanent, apply the following in your terminal:

```
# sysctl -p /etc/sysctl.conf
```

Assumptions

Throughout this configuration example we use the following assumptions for our network infrastructure:

Network Setup

- eth0 is on management network
- eth1 is on inline network

Please refer to the following table for IP and subnet information :

Network Name	Card	Subnet	Gateway	PacketFence Address
eth0	Management	192.168.1.0/24	192.168.1.1	192.168.1.5
eth1	Inline	192.168.2.0/24	192.168.2.1	192.168.2.1

DHCP/DNS

- PacketFence provides its own DHCP service. It will take care of IP address distribution in our Inline network. PacketFence will not provide DHCP services on the management network - this is the responsibility of your own infrastructure.
- PacketFence provides its own DNS service. However for the Inline version, we need to provide the DNS of your infrastructure.

Installation

Import the virtual machine

PacketFence ZEN 5.3.0 comes in a pre-built virtual disk (OVA). If you are using an ESX-type hypervisor, you need to import the OVA using vSphere Client (or vCenter). We are not supporting any Xen-based hypervisors yet.

Import to ESX

Make sure that there is two virtual network cards created. Assign the first card (eth0) to your management network (Production) and assign the second one (eth1) to the Inline network.

Virtual Machine passwords

Management (SSH/Console)

- Login: root
- Password: [p@ck3tf3nc3](#)

Captive Portal Registration User

- Login: demouser
- Password: demouser

Configuration

Inline enforcement configuration

The inline enforcement is a very convenient method for performing access control on older network equipment that is not capable of doing VLAN enforcement or that is not supported by PacketFence.

An important configuration parameter to have in mind when configuring inline enforcement is that the DNS reached by these users should be your actual production DNS server - which shouldn't be in the same broadcast domain as your inline users. The next section shows you how to configure the proper inline interface and it is in this section that you should refer to the proper production DNS.

Inline enforcement uses **ipset** to mark nodes as registered, unregistered and isolated. It is also now possible to use multiple inline interfaces. A node registered on the first inline interface is marked with an IP:MAC tuple (for L2, only ip for L3), so when the node tries to register on an other inline interface, PacketFence detects that the node is already registered on the first inline network. It is also possible to enable `inline.should_reauth_on_vlan_change` to force users to reauthenticate when they change inline network.

By default the inline traffic is forwarded through the management network interface but it is possible to specify another one by adding in `pf.conf` the option `interfaceSNAT` in inline section. It is a comma delimited list of network interfaces like `eth0,eth1.2`. It's also possible to specify a network that will be routed instead of using NAT by adding in `conf/networks.conf` an option `nat=no` under one or more network sections (take care of the routing table of the PacketFence server).

Another important setting is the **gateway** statement. Since it is the only way to get the PacketFence server inline interface IP address, it is mandatory to set it to this IP (which is supposed to be the same as in the `ip` statement of the inline interface in `conf/pf.conf`).

Configuring your PacketFence environment

Before booting your VM, make sure that the PC is correctly connected in the management network and that the link is up.

Once powered, open a browser and point it to the configuration URL as stated by the VM login prompt (ie. https://PF_IP:1443/configurator). The configuration process is a five steps process at the end of which, the VM will be a persistent working PacketFence environment. If you cannot reach the configurator, but you made sure that the connectivity is fine (ie. `ping @PF_IP`) and PF `httpd.admin` is running (ie. `ps-edf | grep httpd.admin` in your terminal or `service packetfence status`), then try to disable iptables.

Step 1: Enforcement

The first and most important step of the configuration process. This is where you'll choose the enforcement technique; either VLAN (out-of-band), INLINE (in-band) or both of them.

The choice(s) made on this step will influence the next step where you'll need to configure the different networks.

Each enforcement mode has its own required interface types that you'll have to configure on step 2.

In this guide we will show you how to configure the INLINE (in-band) mode. If you want to configure the VLAN (out-of-band) mode please refer to the guide [PacketFence Out of Band Deployment Quick Guide ZEN](#).

Step 2: Networks

This step will ask you to statically configure your network interfaces (note that DHCP interfaces configuration is not supported).

The web interface will list all currently installed network interfaces on the system. An IP and a netmask will be visible if the network interface is configured (either by DHCP or already manually configured). You can edit those ones, create/delete VLANs on physical interfaces and enable/disable an interface. Note that these changes are effective on the moment you make them. Persistence will be written only for ENABLED interfaces. Which means that if you change your management ip address, to pursue the configurator, you will need to go on this new ip address you setup.

In all time, you'll need to set a Management interface.

Required interface types for inline enforcement:

```
Management
Inline layer 2
```

Note that you can only set ONE (1) management interface. This one will work for both in the case you choose both modes.

In our case, we will use two interfaces, one will be for the management, the other one will be for the inline.

For the Inline interface, we connect this network card in a switch port in the Inline network.

Here's a sample configuration for both of them:

```
eth0: Management
IP Address: 192.168.1.5
Netmask: 255.255.255.0
Gateway: 192.168.1.254
```

```
eth1: Inline Layer 2
IP Address: 192.168.2.1
Netmask: 255.255.255.0
DNS Servers: 192.168.1.10
```

This configuration take into account that you have an available machine in the management network to access the admin interface of PacketFence.

Make sure that those 2 interfaces are in an Enabled state for the persistence to occur.

We also need to set the Default Gateway which will generally be the gateway of the management network.

Note that if you have a routed network that needs to be taken into account by PacketFence as an Inline subnet, you will need to add this via **Routed Network** and select Inline L3. This configuration is available only from the PacketFence web administrative GUI and not from the configurator.

Once everything's set, click Continue to proceed with the next step.

Step 3: Database Configuration

This step will configure the MySQL server needed by PacketFence. Database and schema will be created as well as the necessary user for operations. Root account will also be secured if necessary (set a password and disallow remote login).

Start the MySQLd service if it is not started. Click the MySQL Start button at the top of the web page Warning! MySQL server does not seems to be running. You should start it to avoid any problems. Start MySQL.

Then you will need to create the root password for MySQL database. Click on the Test button and write a complex password (recommended) twice and save. When you are done creating the password, put the new root password and click on Test to validate it. You should see Success! Successfully connected to the database mysql with user root

Next section will create the database and load the correct schema on it. Simply leave the default database name and click Create databases and tables.

The last section of this step is the PacketFence user account on the MySQL server. Simply leave the default pf username here and choose of a password. This one will automatically be set in the PacketFence configuration where you'll be able to retrieve it in any case. Once the password entered twice, click Create user.

If you got a Success! message for this all three sections, click Continue.

Step 4: PacketFence Configuration

This step will configure the general options of your PacketFence installation. Theses are needed configurations that will most of the time fits customer specifications.

Almost all of the required information here are self-explanatory. The only one that could be confusing is the DHCP Servers section. In this one, enter a comma-delimited list of all the DHCP Server on the customer network so when PacketFence will see DHCP traffic originating from these IPs, no rogue-DHCP alerts will be triggered.

Packetfence will use the domain and the hostname to generate the URL to redirect devices on the captive portal. If you have a HTTP certificate use the same hostname and domain n ame to validate the SSL connection on the captive portal.

In the last section, Local Database Passwords, you will have to chose the password encryption for local accounts (guest automatically generated and manually created account).

Click Continue once all the fields are completed.

Step 5: Administration

This is the step where we create the administrative user to access the PacketFence Administration Web Interface.

Simply provide the desired username and password, then click Create user.

Step 6: Services - Confirmation

The last but not the least. Here, we start the PacketFence server according to the configurations made in the previous steps. If everything goes as expected, you'll be prompted by a window inviting you to continue to the web administration interface.

You'll be able to login to the PacketFence web administration interface with the credentials created in Step 5.

Services status will help you monitor if everything goes as expected. If not, you'll see which service is in trouble and the log output will help you determine the problem that occurs.

PacketFence configuration files

If you want to customize the configuration, we suggest that you take a look into the PacketFence Administration Guide prior doing so. The main configuration files are :

- `conf/pf.conf` : Configuration for the PacketFence services
- `conf/networks.conf` : Definition of the registration and isolation networks to build DNS and DHCP configurations. In our case, we included the registration and isolation networks.

In standard inline enforcement setup, you should not have to modify any configuration file to make things work. Every modification of the configuration is now done only via the admin interface, we DO NOT advise customers to edit the configuration files.

Network Devices

In an inline configuration, the required configurations for network devices (desktops, tablets, printers, etc.) will be to make sure they can all communicate with PacketFence. In other words for a switch you will need to configure every ports on which devices will be connected using the access mode with all of them in the same inline network. Access point will be connected as device to be in the inline subnetwork.

Exemple with a Cisco switch:

You should be in mode `#conf-t` if not execute *configuration terminal* in your CLI.

```
# interface range [port-range]
# switchport mode access vlan 1
# no shutdown
# interface [packetfence_eth1]
# switchport mode access vlan 1
# no shutdown
# end
# copy running-configuration startup-configuration
```

Now you can connect any devices that you want to be in the inline network in any of the port you have just configured.

FreeRADIUS

PacketFence ZEN 5.3.0 comes with a pre-configured FreeRADIUS to do Wired and Wireless 802.1X with EAP as well as MAC Authentication. We created a local user for the 802.1X authentication.

Since we do a Inline mode, we will not use radius.

Inline Access

- Make sure that the Inline and management card are properly configured
- connect a device with a DHCP IP in the Inline subnetwork
- make sure the device are able to communicate with PacketFence on the Inline network card and cannot access the management network

OMAPI

Configuring the DHCP OMAPI (optional)

In order to speed up the IP address lease lookup, you can configure the DHCP OMAPI so that queries for IP and MAC associations are made faster.

First, execute the following command in an SSH session.

```
# dd if=/dev/urandom bs=16 count=1 2>/dev/null | openssl enc -e -base64
```

This should produce an output similar to this :

```
m4NMk0Kc9Ixfwk8cL2fP4g==
```

Now paste the output in the Administration interface under *Configuration/OMAPI/OMAPI base64 key* and save.

The screenshot shows the PacketFence Administration interface. The top navigation bar includes 'Status', 'Reports', 'Nodes', 'Users', and 'Configuration'. The left sidebar lists various configuration categories, with 'OMAPI' selected. The main content area is titled 'OMAPI' and contains the following configuration options:

- Enabled ip2mac lookup using OMAPI** (checked): Use OMAPI to query DHCPd for the MAC address of a given IP address
- Enabled mac2ip lookup using OMAPI** (checked): Use OMAPI to query DHCPd for the IP address of a given MAC address
- OMAPI Key name**: pf_omapi_key (The OMAPI key name for signing messages)
- OMAPI base64 key**: MI8L5G+c+SPr0zYegrnxLg== (The OMAPI base64 key for signing messages)
- OMAPI Port**: 7911 (The OMAPI port number)
- OMAPI host**: localhost (The OMAPI host)

At the bottom of the configuration area, there are 'Save' and 'Reset' buttons.

Now restart the dhcpd service using the following command in an SSH session.

```
# /usr/local/pf/bin/pfcmd service dhcpd restart
```

Traffic shaping

Since PacketFence 5.2 it's now possible to shape the inline traffic based on the role of the device.

How we classify

If you launch:

```
# ipset -L
Name: PF-iL2_ID1_192.168.2.0
Type: bitmap:ip
Revision: 0
Header: range 192.168.2.0-192.168.2.255
Size in memory: 152
References: 2
Members:

Name: PF-iL2_ID2_192.168.2.0
Type: bitmap:ip
Revision: 0
Header: range 192.168.2.0-192.168.2.255
Size in memory: 152
References: 2
Members:

Name: PF-iL2_ID3_192.168.2.0
Type: bitmap:ip
Revision: 0
Header: range 192.168.2.0-192.168.2.255
Size in memory: 152
References: 2
Members:
```

You can see that PacketFence created 3 new ipset sessions based on the inline network ip and on the role id defined in Roles section (Configuration → Roles, to see the id of each role).

So when a device will register on the captive portal, PacketFence will add the device in the corresponding ipset session (role id, network).

Next iptables rules in mangle table will classify the traffic based on the ipset session:

```
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID1_192.168.2.0 src -j
CLASSIFY --set-class 1:1
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID1_192.168.2.0 dst -j
CLASSIFY --set-class 1:1
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID2_192.168.2.0 src -j
CLASSIFY --set-class 1:2
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID2_192.168.2.0 dst -j
CLASSIFY --set-class 1:2
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID3_192.168.2.0 src -j
CLASSIFY --set-class 1:3
-A postrouting-int-inline-if -m set --match-set PF-iL2_ID3_192.168.2.0 dst -j
CLASSIFY --set-class 1:3
```

So here the role id 1 will have the class 1:1.

Configure Traffic shaping

Here 2 examples of tc rules, the first one will apply an upload/download of: 1mb/1mb on role id 1 2mb/2mb on role id 2 3mb/3mb on role id 3

```

tc qdisc del dev eth0 root
tc qdisc add dev eth0 root handle 1:0 htb default 1

tc class add dev eth0 parent 1:0 classid 1:1 htb rate 1mbit ceil 1mbit
tc class add dev eth0 parent 1:0 classid 1:2 htb rate 2mbit ceil 2mbit
tc class add dev eth0 parent 1:0 classid 1:3 htb rate 3mbit ceil 3mbit
tc qdisc add dev eth0 parent 1:1 sfq
tc qdisc add dev eth0 parent 1:2 sfq
tc qdisc add dev eth0 parent 1:3 sfq

tc qdisc del dev eth1 root
tc qdisc add dev eth1 root handle 1:0 htb default 1

tc class add dev eth1 parent 1:0 classid 1:1 htb rate 1mbit ceil 1mbit
tc class add dev eth1 parent 1:0 classid 1:2 htb rate 2mbit ceil 2mbit
tc class add dev eth1 parent 1:0 classid 1:3 htb rate 3mbit ceil 3mbit
tc qdisc add dev eth1 parent 1:1 sfq
tc qdisc add dev eth1 parent 1:2 sfq
tc qdisc add dev eth1 parent 1:3 sfq

```

The second one will apply an upload/download of: 1mb/10mb on role id 1 2mb/20mb on role id 2 3mb/30mb on role id 3

```

tc qdisc del dev eth0 root
tc qdisc add dev eth0 root handle 1:0 htb default 1

tc class add dev eth0 parent 1:0 classid 1:1 htb rate 1mbit ceil 1mbit
tc class add dev eth0 parent 1:0 classid 1:2 htb rate 2mbit ceil 2mbit
tc class add dev eth0 parent 1:0 classid 1:3 htb rate 3mbit ceil 3mbit
tc qdisc add dev eth0 parent 1:1 sfq
tc qdisc add dev eth0 parent 1:2 sfq
tc qdisc add dev eth0 parent 1:3 sfq

tc qdisc del dev eth1 root
tc qdisc add dev eth1 root handle 1:0 htb default 1

tc class add dev eth1 parent 1:0 classid 1:1 htb rate 10mbit ceil 10mbit
tc class add dev eth1 parent 1:0 classid 1:2 htb rate 20mbit ceil 20mbit
tc class add dev eth1 parent 1:0 classid 1:3 htb rate 30mbit ceil 30mbit
tc qdisc add dev eth1 parent 1:1 sfq
tc qdisc add dev eth1 parent 1:2 sfq
tc qdisc add dev eth1 parent 1:3 sfq

```


Test

Register a device in inline enforcement

You can now test the registration process. In order to do so:

- connect an unregistered device into the switch
- make sure PacketFence provides an IP address to the device. Look into the following log file: `/usr/local/pf/logs/packetfence.log` or verify on the computer you obtain an ip the right subnet range

From the computer:

- open a web browser
- try to connect to a HTTP site (Not HTTPS, eg. <http://www.google.com>)
- make sure that whatever site you want to connect to, you have only access to the registration page.

Register the computer using the following informations:

- user: demouser
- password: demouser

Once a computer has been registered:

- make sure PacketFence changes the firewall (ipset -L) rules so that the user is authorized through. Look into PacketFence log file: `/usr/local/pf/logs/packetfence.log`
- from the web administrative interface, go under Nodes and make sure you see the computer as *Registered*.
- the computer has access to the network and the Internet.

Additional Information

For more information, please consult the mailing archives or post your questions to it. For details, see:

- packetfence-announce@lists.sourceforge.net: Public announcements (new releases, security warnings etc.) regarding PacketFence
- packetfence-devel@lists.sourceforge.net: Discussion of PacketFence development
- packetfence-users@lists.sourceforge.net: User and usage discussions

Commercial Support and Contact Information

For any questions or comments, do not hesitate to contact us by writing an email to: support@inverse.ca.

Inverse (<http://inverse.ca>) offers professional services around PacketFence to help organizations deploy the solution, customize, migrate versions or from another system, performance tuning or aligning with best practices.

Hourly rates or support packages are offered to best suit your needs.

Please visit <http://inverse.ca/> for details.

GNU Free Documentation License

Please refer to <http://www.gnu.org/licenses/fdl-1.2.txt> for the full license.