Product Manual

XG-7100-1U

Netgate

Sep 21, 2018

CONTENTS

Thank you for your purchase of the pfSense® XG-7100 1U System. This Netgate appliance provides a powerful, reliable, cost-effective solution.

Quick Start Guide

The Quick Start Guide covers the first time connection procedures and will provide you with the information you need to get your appliance up and running.

CHAPTER

ONE

I/O PORTS

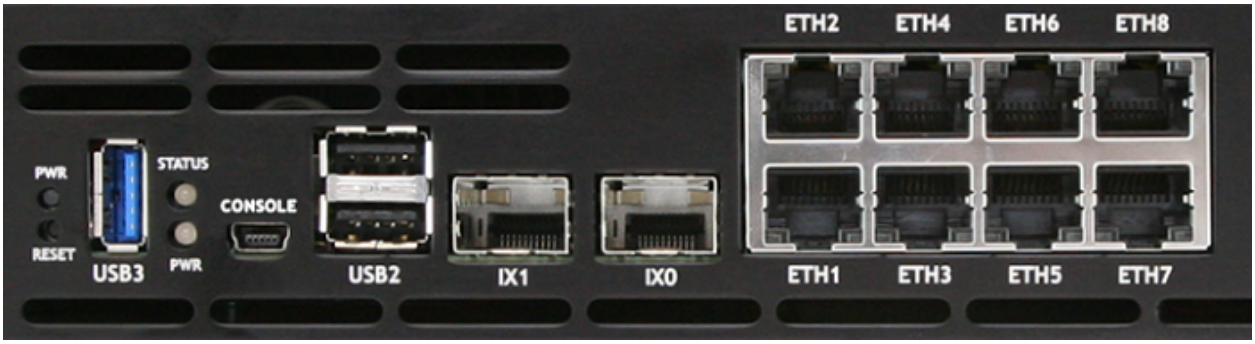

Ports are assigned as pictured.

1.1 Ethernet Ports

Interface Name	Port Name	Port Type	Port Speed
WAN	ETH1	RJ-45	1 Gbps
LAN	ETH2-ETH8	RJ-45	1 Gbps
OPT1	IX0	SFP+	10 Gbps
OPT2	IX1	SFP+	10 Gbps

Note: ETH1-8 are **switched ports** sharing 5 Gbps (2x 2.5 Gbps) to the Intel SoC. These ports can be isolated as an independent interface with the configuration of VLAN tagging as shown in *XG-7100 Switch Overview*.

1.2 Optional 4-Port Intel 1 Gb Ethernet Expansion Card

Ports are assigned as numbered.

Num	Interface Name	Port Name	Port Type	Port Speed
0	OPT3	igb0	RJ-45	1 Gbps
1	OPT4	igb1	RJ-45	1 Gbps
2	OPT5	igb2	RJ-45	1 Gbps
3	OPT6	igb3	RJ-45	1 Gbps

Warning: High Availability (HA) can be used, but currently there is one restriction when it comes to configuring switchports for HA. Because the uplinks from the switch to the SoC are always up, failover is only effective in scenarios where a system completely dies. If a single switch interface goes down, CARP will not be able to detect this properly so the **PRIMARY** will remain **PRIMARY** on any switch interfaces that drop link.

The **SECONDARY** will also consider itself **PRIMARY** of the network associated to the switch link that dropped. In this situation, LAN clients will likely go through the **SECONDARY** but will not be able to get online if NAT is utilized with a WAN CARP IP. It's possible to NAT to the WAN interface IP to get around this but it can cause state issues during failover.

There is an Intel-supplied driver issue, which is noted in the Intel Release Notes for the C3000, preventing 1Gbps and 10Gbps copper modules from being recognized on the SFP+ ports. Copper modules are not supported.

LAGG is not currently supported on the ethernet switchports.

This will be addressed in a future pfSense release.

1.3 Other Ports, Buttons, and Indicators

- Semi-recessed Power (PWR) (performs a graceful shutdown of pfSense software)
- Recessed Reset Button (performs a hard reset, immediately turning the system off)
- 1x USB 3.0
- Status LED
- Power (PWR) LED (green when powered on, red after a graceful shutdown)
- Console (Mini-USB)
- 2x USB 2.0

Note: When a graceful shutdown is performed, the XG-7100 Power (PWR) LED will turn red but will stay lit. The Ethernet activity LEDs will turn off. The power supply fan will continue to run. Turning off the rocker switch on the back of the power supply will eliminate all power to the system.

The power button should be depressed 3-5 seconds to initiate a graceful shutdown or to power on the device when the PWR LED is red.

Warning: A hard reset of the system could cause data corruption and should be avoided. Halt or reboot the system through the console menu or the web configurator to avoid data corruption.

CHAPTER

TWO

XG-7100 SWITCH OVERVIEW

2.1 Interface Links

In addition to two SFP+ interfaces, there is also an ethernet switch on the XG-7100. There are eight ethernet ports on this switch that are physically accessible - these interfaces are referred to as ETH1-ETH8. In addition to those 8 ports, there are also three additional ports that operate behind the scenes - PORT 0, PORT 9 (ix2), and PORT 10 (ix3).

ETH1-ETH8 are gigabit switchports.

PORT 9-10 are 2.5 Gbps uplink switchports. These two ports connect the ethernet switch to a Denverton SoC. The SFP+ interfaces (ix0 and ix1) also connect to this SoC.

The diagram below demonstrates how these interfaces are connected:

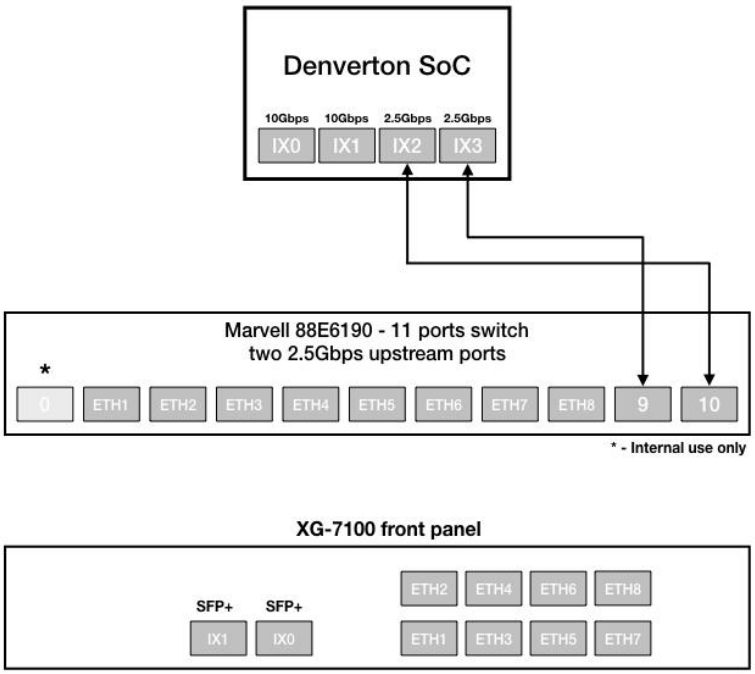

From the operating systems perspective, there are four physical interfaces present:

```
ix0 - 10Gbps SFP+
ix1 - 10Gbps SFP+
ix2 - 2.5 Gbps (2500-Base-KX, switch link to SoC/CPU)
ix3 - 2.5 Gbps (2500-Base-KX, switch link to SoC/CPU)
```

2.2 Switch LAGG

ix2 and ix3 (switch uplink ports 9 and 10), are configured as a load-balanced LAGG. This provides an aggregate uplink capable of 5Gbps for ethernet switchports ETH1-8. This is further demonstrated in the diagram below:

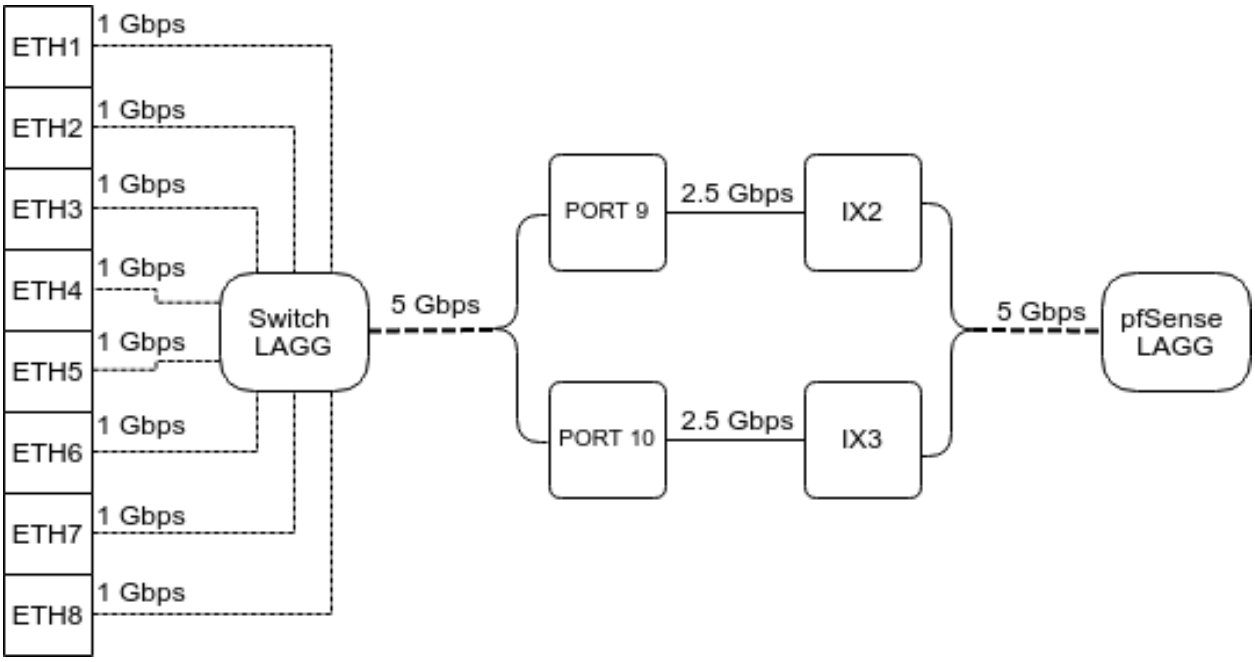

When data is received on ETH1-8, the switch is capable of utilizing LAGG to determine whether that data should be sent out of PORT 9 or PORT 10. That data then passes over one of two 2.5Gbps switch links (PORT 9/10) to the SoC. Data coming from PORT 9 has a direct line to ix2 and data from PORT 10 has a direct line to ix3.

pfSense LAGG will then take in traffic from both **ix2** and **ix3** as though it came in on a single interface, **lagg0**. The same concept applies to traffic sourcing from the pfSense LAGG to the switch LAGG.

2.3 Switch VLANs

By default, ETH1 on the the switch is configured as a WAN interface and ETH2-8 are configured as the LAN interface. These eight switchports are customizable and each can be configured to act as an independent interface. For example, all of these configurations are possible:

- ETH1-8 dedicated as a LAN switch
- ETH1-4 configured as a switch for LAN network A and ETH5-8 configured as a switch for LAN network B
- ETH1-8 configured as individual network interfaces
- ETH1 configured for WAN A, ETH2 configured for WAN B, ETH3 configured for LAN network A, ETH4-6 configured as a switch for LAN network B, and ETH8 configured as a H/A sync port.

These scenarios are possible by utilizing VLANs. Each of the switchports (ETH1-8 and PORT9-10) are VLAN aware interfaces. They are capable of functioning like a standard access or trunk port:

Access Port: Adds a VLAN tag to inbound untagged traffic

Trunk Port: Allows tagged traffic containing specified VLAN IDs

In the default configuration, two VLANs are used to create the ETH1 WAN interface and ETH2-8 LAN interface:

WAN	VLAN 4090
LAN	VLAN 4091

ETH1-8 are configured to act as **Access** ports.

- When data comes into the ETH1 interface, a VLAN tag of 4090 is added to the ethernet frame.
- When data comes into interfaces ETH2-8, a VLAN tag of 4091 is added to the ethernet frame.

PORT9-10 are configured to act as **Trunk** ports.

- By default, only ethernet frames containing a VLAN tag of 4090 or 4091 are allowed over the trunk.

Each VLAN configured on the switch uses the LAGG interface as its parent interface. For example, the default interface assignment for WAN and LAN:

WAN	lagg0.4090
LAN	lagg0.4091

This means **vlan4090** and **vlan4091**, as well as any other VLANs created for the switch, all share the same 5Gbps LAGG uplink across two 2.5Gbps links. The visual below demonstrates how the VLAN tagging works along with the traffic flow:

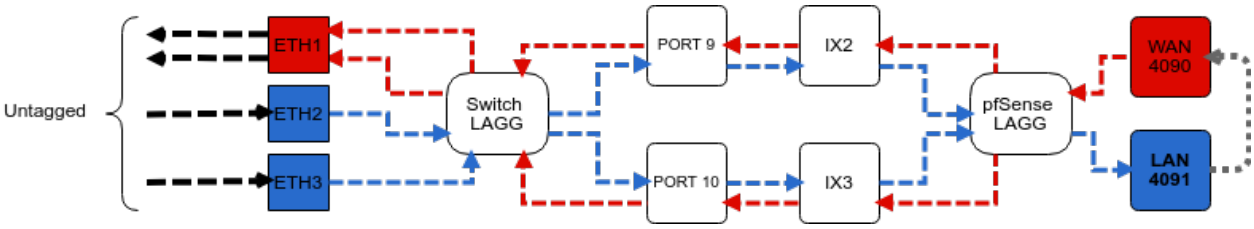

Note that traffic leaving and entering the ETH1-3 interfaces in the visual above are untagged. Devices sending/receiving traffic over these ports do not need to be VLAN aware. The VLAN tagging that occurs within the switch is completely transparent to clients. It's used solely for segmenting switch traffic internally.

Aside from being able to specify whether a switchport should act as an access or trunk port, it's also possible to disable 802.1q VLAN mode. When this is done, a third mode called **Port VLAN Mode** is enabled. In this mode, any and all VLAN tags are allowed on all ports. No VLAN tags are added or removed. Think of it as a dummy switch that retains VLAN tags on frames, if present. This mode is useful when you have numerous VLANs on your network and want to physically segment the switch, while allowing the same VLANs on all segments of the switch.

In **Port VLAN Mode**, rather than specifying which interfaces are associated to a VLAN, you can specify which physical ports form a switch. For example, if I want to create two physical switches that act as individual dummy switches - allowing tagged or untagged traffic, I could configure **Port VLAN Mode** like so:

```
// UPLINKS
VLAN group 9, Port 9, Members 1,2,3,4,10
VLAN group 10, Port 10, Members 1,2,3,4,9

// SWITCH-A
VLAN group 1, Port 1, Members 2,3,4,9,10
VLAN group 2, Port 2, Members 1,3,4,9,10
VLAN group 3, Port 3, Members 1,2,4,9,10
VLAN group 4, Port 4, Members 1,2,3,9,10

// SWITCH-B
VLAN group 5, Port 5, Members 6,7,8
VLAN group 6, Port 6, Members 5,7,8
VLAN group 7, Port 7, Members 5,6,8
VLAN group 8, Port 8, Members 5,6,7
```

With this configuration in place, ETH1-8 now function like so:

```
// SWITCH-A
PORT 1 = ETH1
PORT 2 = ETH2
PORT 3 = ETH3
PORT 4 = ETH4
PORT 9 = UPLINK 1
PORT 10 = UPLINK 2

// SWITCH-B
PORT 5 = ETH5
PORT 6 = ETH6
PORT 7 = ETH7
PORT 8 = ETH8
```

SWITCH-A

ETH1-4 can talk to each other and to the LAGG uplink. PORT9-10 are members of this switch. . . this is required for this switch to have uplink to pfSense.

SWITCH-B

ETH5-8 can talk to each other but because PORT9-10 are not included as members, clients connecting to ETH5-8 can only talk to other clients on ETH5-8. They will not be able to reach the SoC where ix2 and ix3 are defined, so they never reach pfSense. This can be useful if you want a device other than pfSense to act as the primary uplink for those connected clients.

Since WAN and LAN are assigned to **lagg0.4090** and **lagg0.4091**, if **Port VLAN Mode** is enabled, be sure to update the LAN and WAN interface assignment to reference the appropriate VLAN. Also remember to create the new VLANs with **lagg0** as the parent interface.

If **Port VLAN Mode** is being used to handle untagged traffic, the **LAGG0** interface should be added, enabled, and configured under Interface Assignments.

2.4 Configuring the Switch

2.4.1 Switch Section

From the pfSense webGUI, there is a menu option called **Switches** under the Interfaces drop-down. This section contains switch specific configuration options.

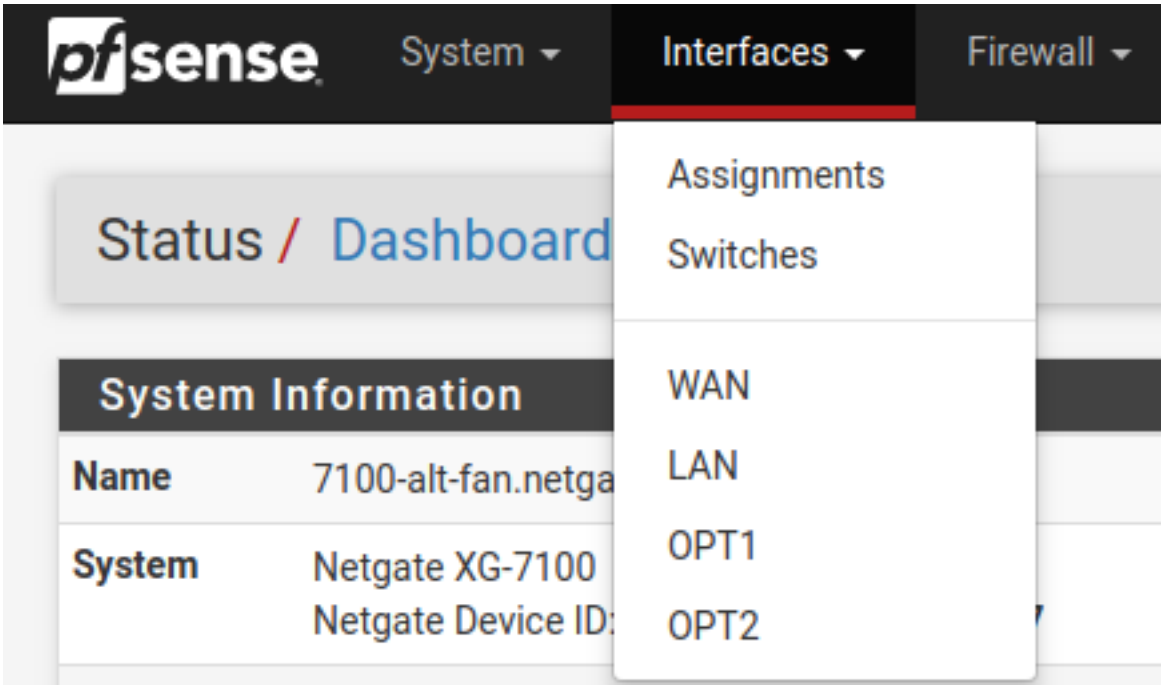

Selecting **Switches** from the drop-down will bring up the Switch page with four sections:

System

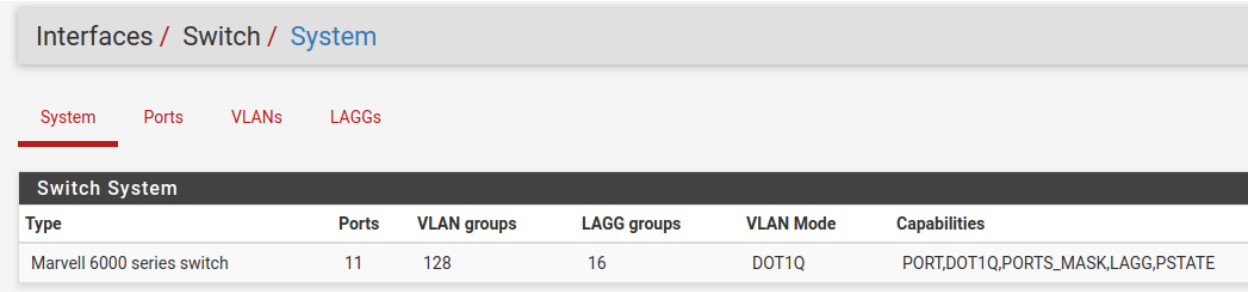

Fig. 1: Information on the Marvell 6000 switch

LAGGs

Ports

Information on switchport status and port names. If **802.1q** is enabled, this section can also be used to specify the native VLAN ID for each port. The Port VID defined will be used to tag inbound untagged traffic.

VLANs

Enable/Disable 802.1q VLAN mode. Configure VLAN access/trunk interfaces with 802.1q or configure port groups with **Port VLAN Mode**.

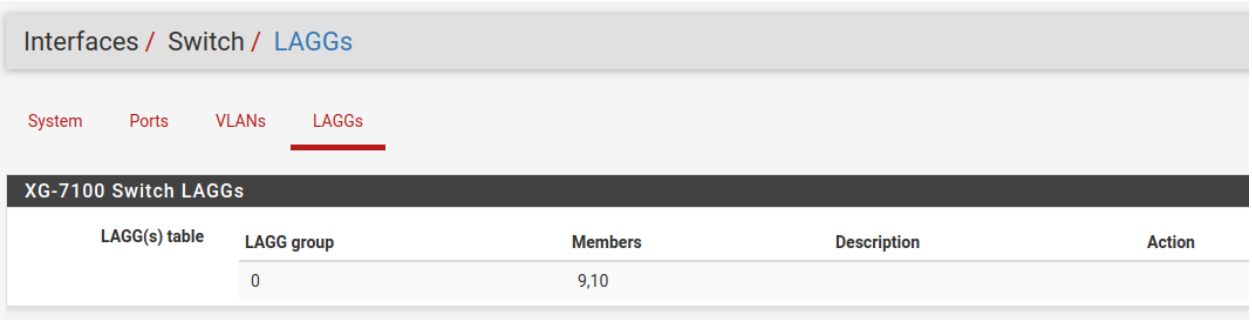

Fig. 2: Information on members of the switch LAG

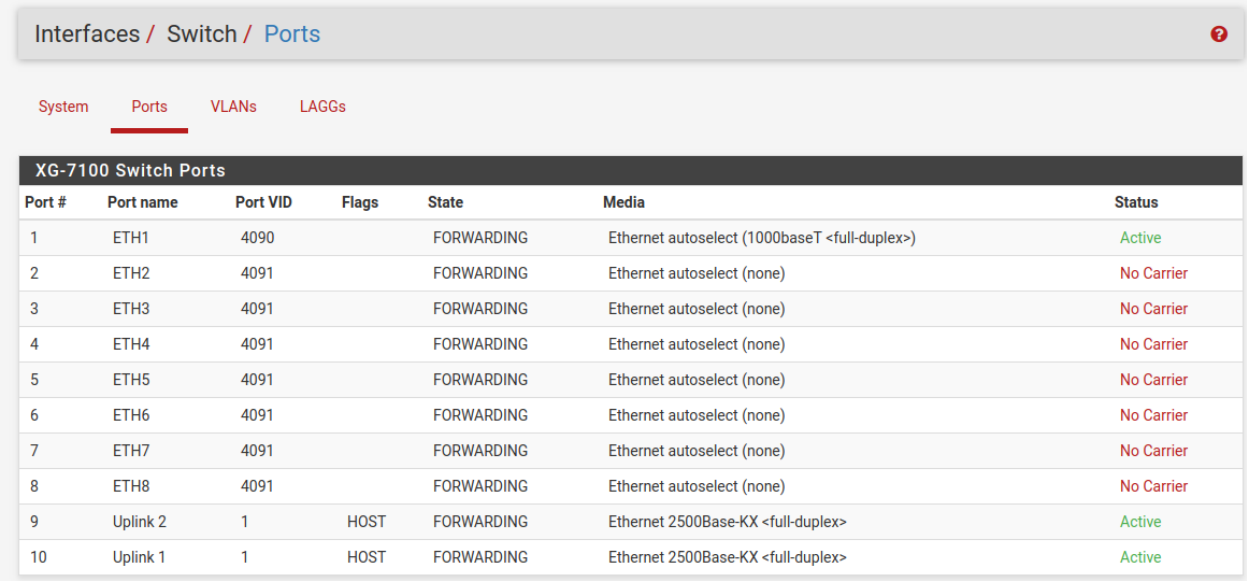

Fig. 3: 802.1q enabled (default)

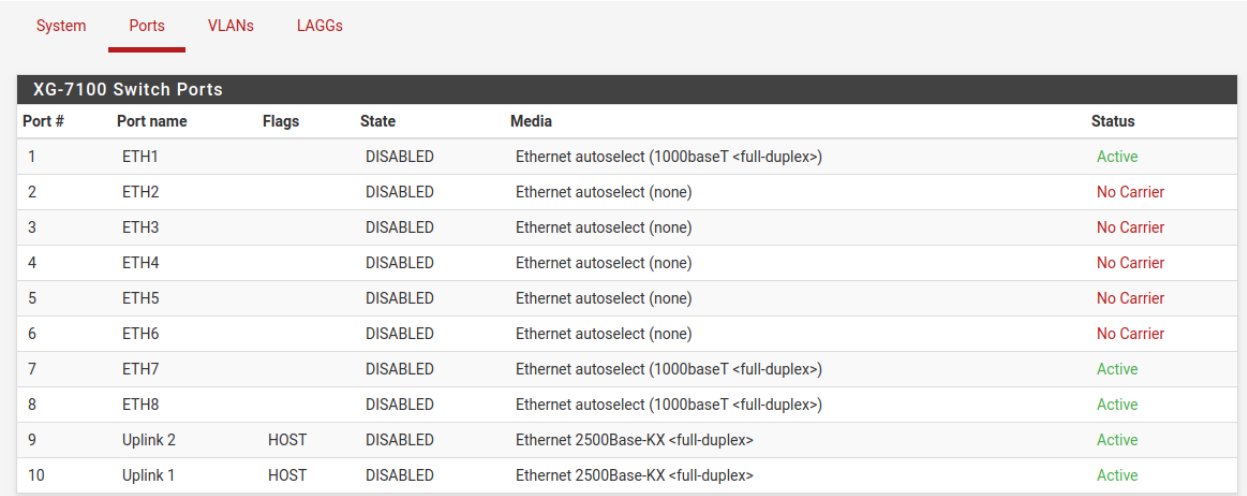

Fig. 4: Port VLAN Mode

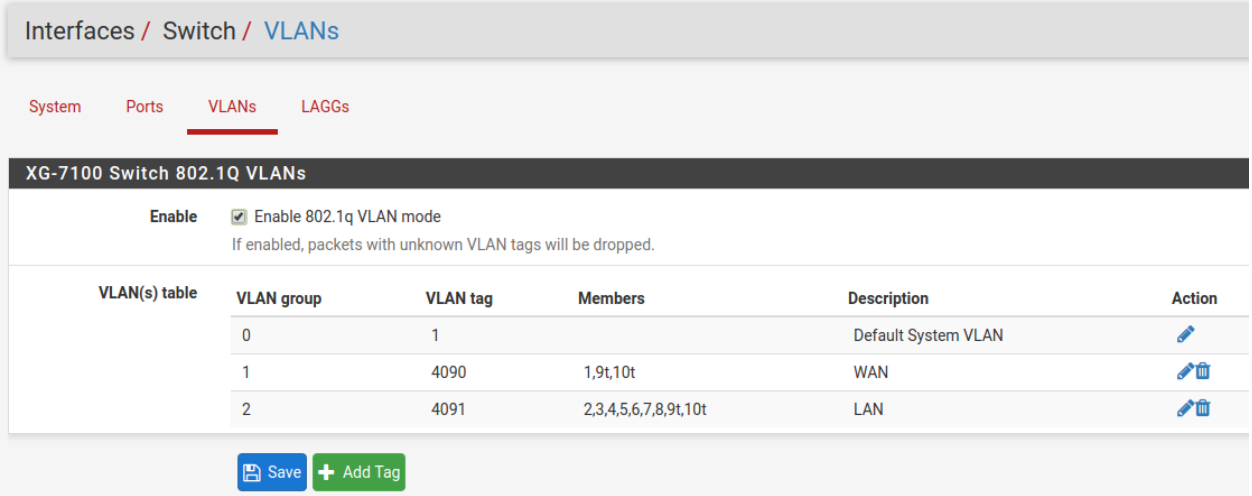

Fig. 5: 802.1q enabled (default)

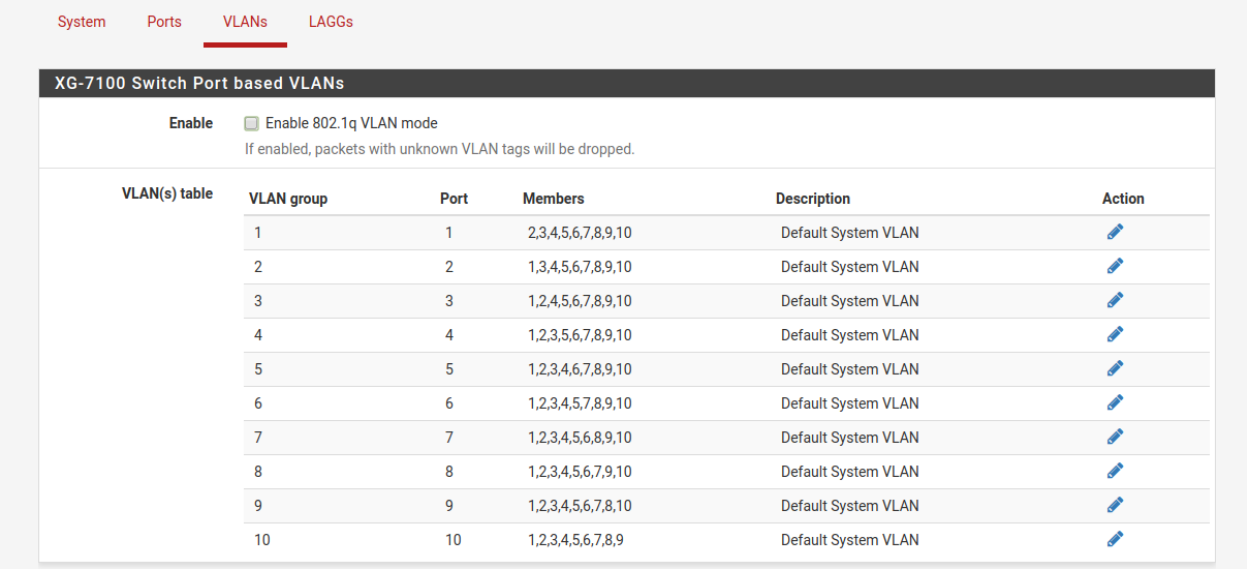

Fig. 6: Port VLAN Mode

2.4.2 Interfaces Section

There is also relevant configurations under **Interfaces -> Assignments**.

Interface Assignments

Under **Interface Assignments**, notice **LAGG0 (UPLINK)** is displayed as an available port but is not enabled in the list of interfaces. This is because the default configuration is only expecting VLAN tagged traffic so the VLAN child interface **4090** and **4091** are enabled instead.

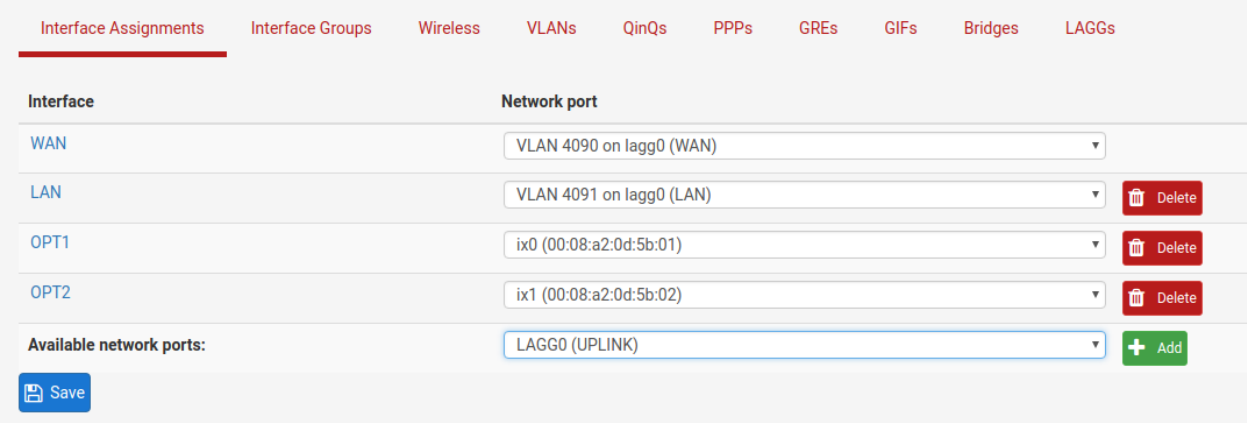

VLANs

Under VLANs, the default WAN and LAN VLAN can be seen. Additional VLAN networks that will be used by the switch should be defined here with **lagg0** as the parent interface.

Any additional VLAN interface added to the switch should also be added, enabled, and configured under Interface Assignments. Firewall rules will also be needed for new interfaces added.

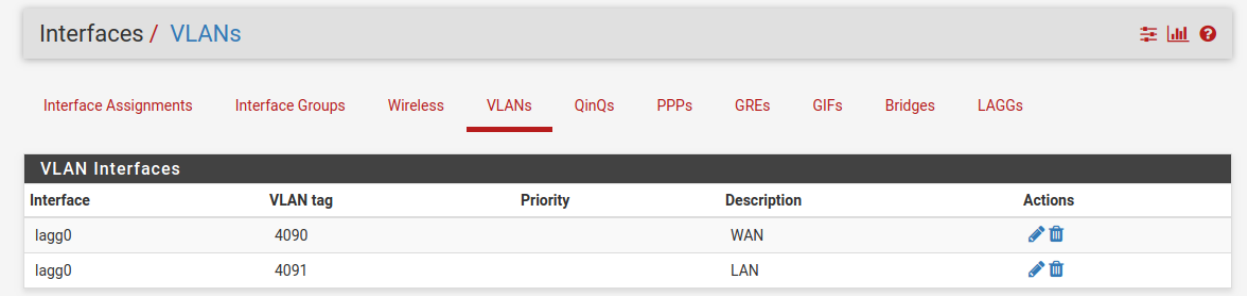

LAGGs

Under LAGGs, the default **lagg0** containing **ix2** and **ix3** can be seen. The **lagg0** interface should not be modified.

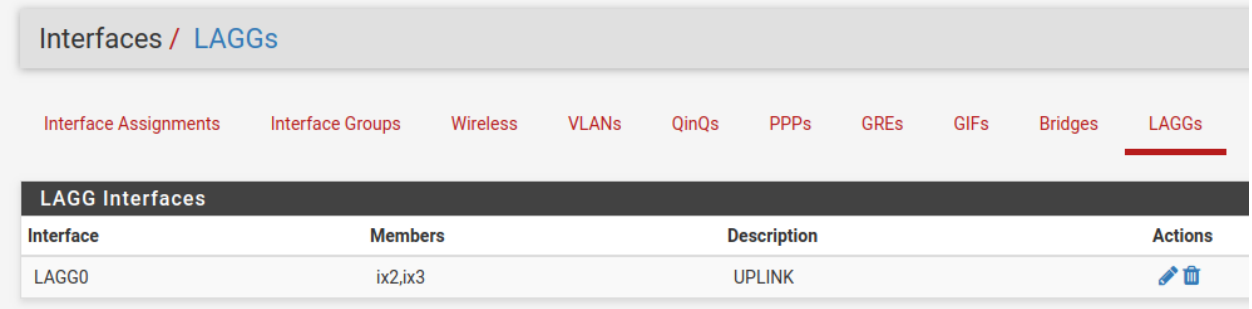

2.5 Switch Configuration Examples

2.5.1 Dedicated LAN switch

In this scenario, SFP+ port ix0 will be configured as the WAN interface. ETH1-8 will be configured as a LAN switch.

For this specific example, I'll perform the WAN interface reassignment over console. Re-assigning the WAN can be done from the webGUI as well.

This is what the default interface assignments look like on a XG-7100 without an addon NIC:

```
*** Welcome to pfSense 2.4.3-RELEASE (amd64) on pfSense ***

 WAN (wan)       -> lagg0.4090 -> v4/DHCP4: 10.10.30.18/24
 LAN (lan)       -> lagg0.4091 -> v4: 192.168.1.1/24
 OPT1 (opt1)     -> ix0        ->
 OPT2 (opt2)     -> ix1        ->

 0) Logout (SSH only)                  9) pfTop
 1) Assign Interfaces                 10) Filter Logs
 2) Set interface(s) IP address       11) Restart webConfigurator
 3) Reset webConfigurator password    12) PHP shell + pfSense tools
 4) Reset to factory defaults         13) Update from console
 5) Reboot system                     14) Enable Secure Shell (sshd)
 6) Halt system                       15) Restore recent configuration
 7) Ping host                         16) Restart PHP-FPM
 8) Shell

Enter an option:
```

In this example, **ix0** will be **WAN**, so select option **1** to re-assign **WAN** from **lagg0.4090** to **ix0**:

```
Enter an option: 1


Valid interfaces are:

ix0     00:a0:c9:00:00:00 (down) Intel(R) PRO/10GbE PCI-Express Network Driver,
ix1     34:12:78:56:01:01 (down) Intel(R) PRO/10GbE PCI-Express Network Driver,
ix2     00:a0:c9:00:00:02 (down) Intel(R) PRO/10GbE PCI-Express Network Driver,
ix3     00:a0:c9:00:00:02 (down) Intel(R) PRO/10GbE PCI-Express Network Driver,

Do VLANs need to be set up first?
If VLANs will not be used, or only for optional interfaces, it is typical to
say no here and use the webConfigurator to configure VLANs later, if required.

Should VLANs be set up now [y|n]?
```

No additional VLANs are needed for this, so enter `n` to continue.

Input **ix0** as the new **WAN** interface name:

```
Should VLANs be set up now [y|n]? n


VLAN interfaces:

lagg0.4090      VLAN tag 4090, parent interface lagg0
lagg0.4091      VLAN tag 4091, parent interface lagg0

If the names of the interfaces are not known, auto-detection can
be used instead. To use auto-detection, please disconnect all
interfaces before pressing 'a' to begin the process.

Enter the WAN interface name or 'a' for auto-detection
(ix0 ix1 ix2 ix3 lagg0.4090 lagg0.4091 or a):
```

Input the same default **LAN** interface of **lagg0.4091** for the **LAN** interface name and press `Enter` to complete the interface reassignment:

```
Enter the WAN interface name or 'a' for auto-detection
(ix0 ix1 ix2 ix3 lagg0.4090 lagg0.4091 or a): ix0

Enter the LAN interface name or 'a' for auto-detection
NOTE: this enables full Firewalling/NAT mode.
(ix1 ix2 ix3 lagg0.4090 lagg0.4091 a or nothing if finished): lagg0.4091

Optional interface 1 description found: OPT1
Enter the Optional 1 interface name or 'a' for auto-detection
(ix1 ix2 ix3 lagg0.4090 a or nothing if finished):

The interfaces will be assigned as follows:

WAN  -> ix0
LAN  -> lagg0.4091

Do you want to proceed [y|n]?
```

The interface assignments should show like this now:

```
*** Welcome to pfSense 2.4.3-RELEASE (amd64) on pfSense ***

 WAN (wan)       -> ix0        -> v4/DHCP4: 10.10.50.111/24
 LAN (lan)       -> lagg0.4091 -> v4: 192.168.1.1/24

 0) Logout (SSH only)                  9) pfTop
 1) Assign Interfaces                 10) Filter Logs
 2) Set interface(s) IP address       11) Restart webConfigurator
 3) Reset webConfigurator password    12) PHP shell + pfSense tools
 4) Reset to factory defaults         13) Update from console
 5) Reboot system                     14) Enable Secure Shell (sshd)
 6) Halt system                       15) Restore recent configuration
 7) Ping host                         16) Restart PHP-FPM
 8) Shell

Enter an option:
```

At this point SFP+ port **ix0** is now configured as the WAN interface. The LAN interface is still configured the same as the default. Next, the switch will need to be updated so that ETH1 (previously WAN) acts the same as ETH2-8. This will be done from the webGUI.

From the webGUI, pull up the Switch VLAN configuration under **Interfaces -> Switches -> VLANs**:

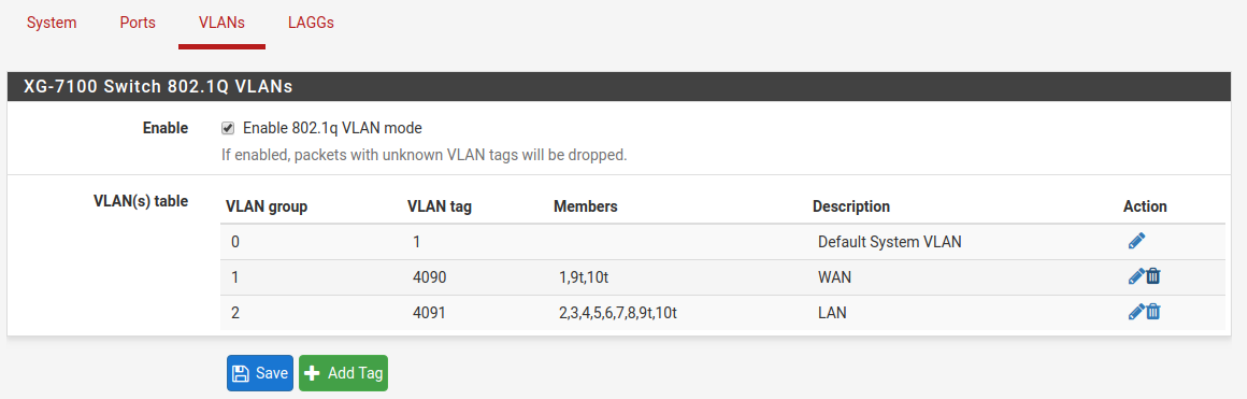

VLAN 4090 is no longer needed since **WAN** is dedicated to **ix0** now. You can either select on the row containing **4090** to delete this entry, or click to remove port 1 as a member:

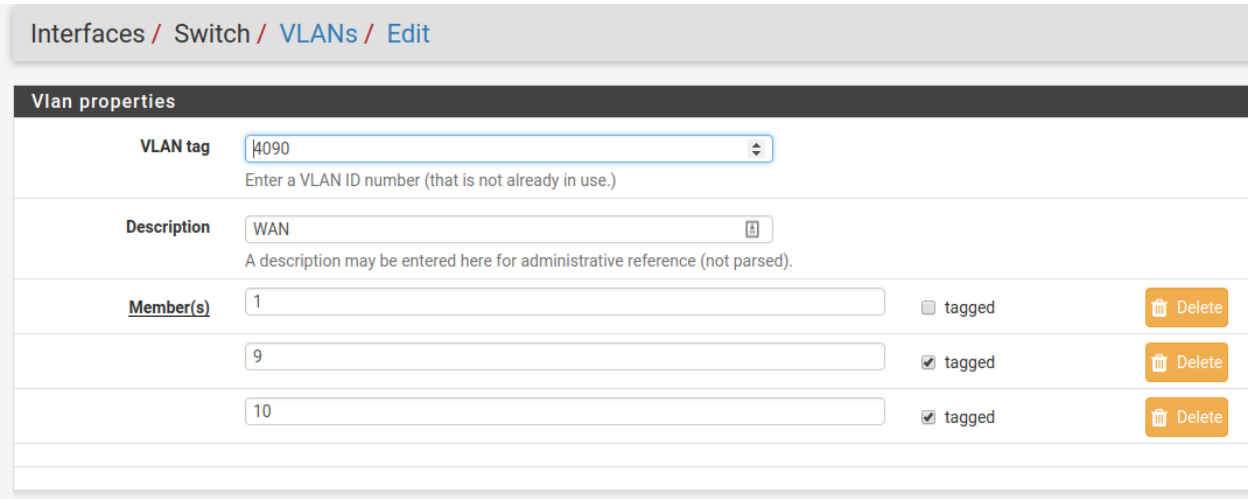

For this example, I simply removed VLAN 4090 from the switch with . Now edit the VLAN 4091 entry to include Member 1 as shown below:

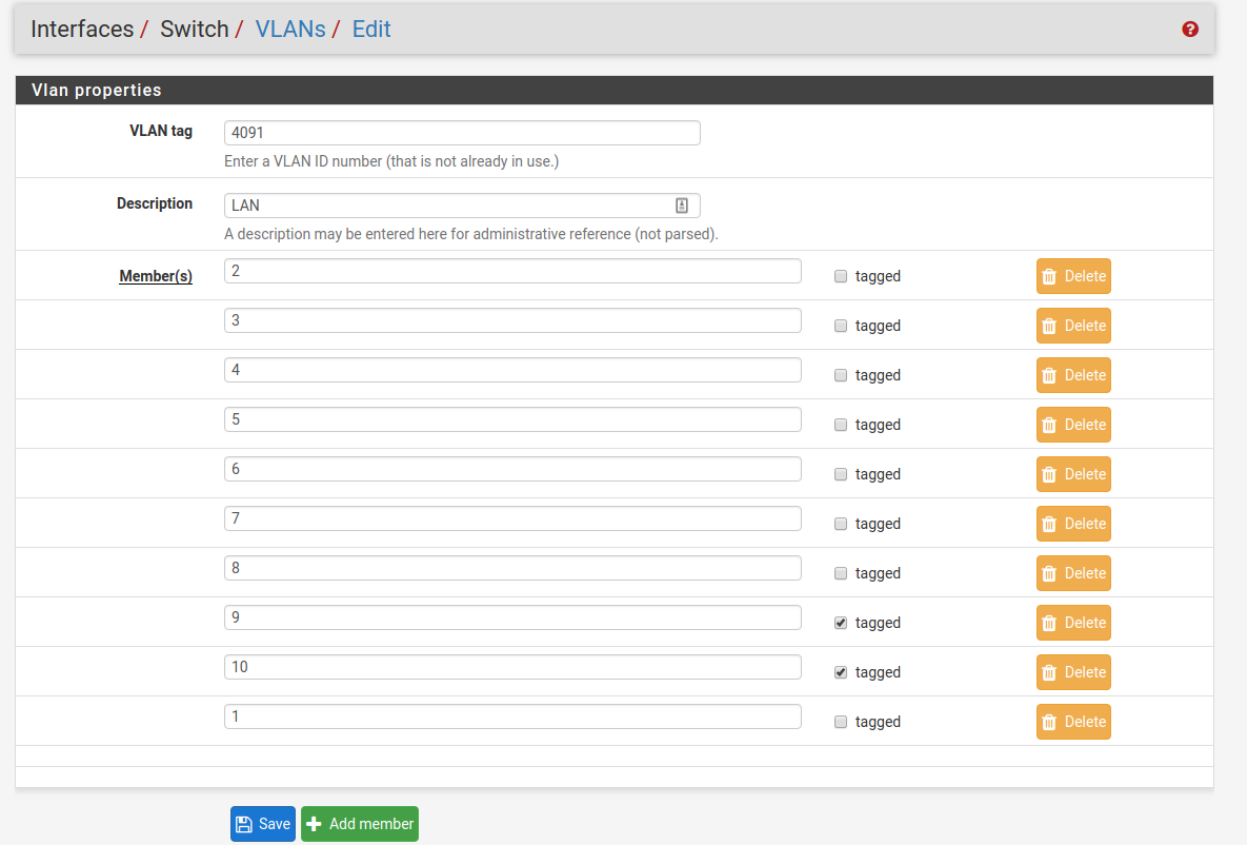

Next, update the PVID for ETH1 so that it uses VLAN 4091 rather than the old VLAN 4090. To do this, click on the **Ports** tab and click on the 4090 Port VID to modify it:

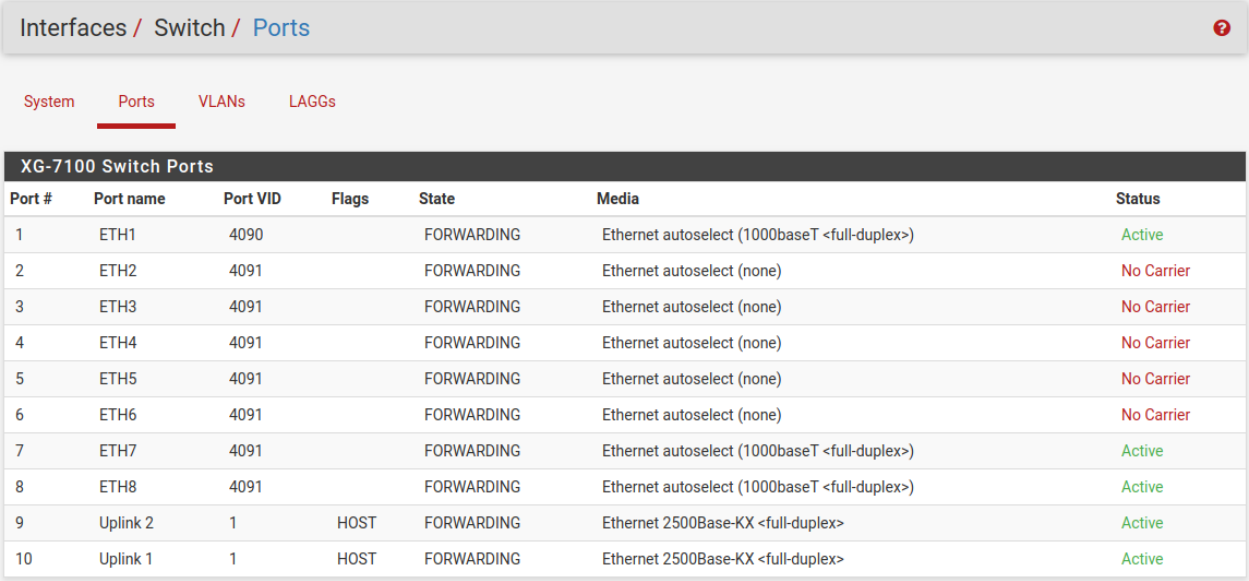

Port #	Port name	Port VID	Flags	State	Media	Status
1	ETH1	4090		FORWARDING	Ethernet autoselect (1000baseT <full-duplex>)	Active
2	ETH2	4091		FORWARDING	Ethernet autoselect (none)	No Carrier
3	ETH3	4091		FORWARDING	Ethernet autoselect (none)	No Carrier
4	ETH4	4091		FORWARDING	Ethernet autoselect (none)	No Carrier
5	ETH5	4091		FORWARDING	Ethernet autoselect (none)	No Carrier
6	ETH6	4091		FORWARDING	Ethernet autoselect (none)	No Carrier
7	ETH7	4091		FORWARDING	Ethernet autoselect (1000baseT <full-duplex>)	Active
8	ETH8	4091		FORWARDING	Ethernet autoselect (1000baseT <full-duplex>)	Active
9	Uplink 2	1	HOST	FORWARDING	Ethernet 2500Base-KX <full-duplex>	Active
10	Uplink 1	1	HOST	FORWARDING	Ethernet 2500Base-KX <full-duplex>	Active

Then click on **Save**:

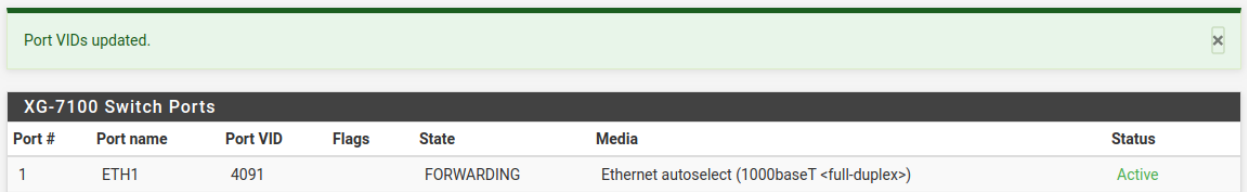

At this point, everything should be configured properly. ETH1-8 will act as a single LAN switch. One final step that should be performed is to remove the old VLAN 4090 from pfSense. So far VLAN 4090 was only removed from the switch. To remove the old VLAN, go to **Interfaces -> Assignments -> VLANs** and use on the 4090 row to remove this VLAN interface:

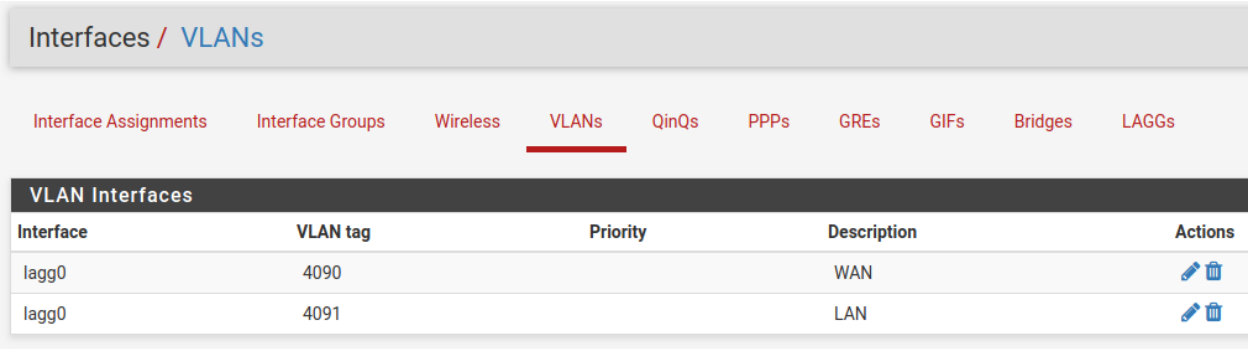

2.5.2 Two LAN switches

In this scenario, the LAN switch from the previous example will be split into two LAN switches.

A new LAN network should be created in pfSense first. Similar to the existing LAN interface, another VLAN interface should be used so the switch can segment traffic appropriately.

Create a new VLAN with **lagg0** as the parent under **Interfaces -> Assignments -> VLANs**:

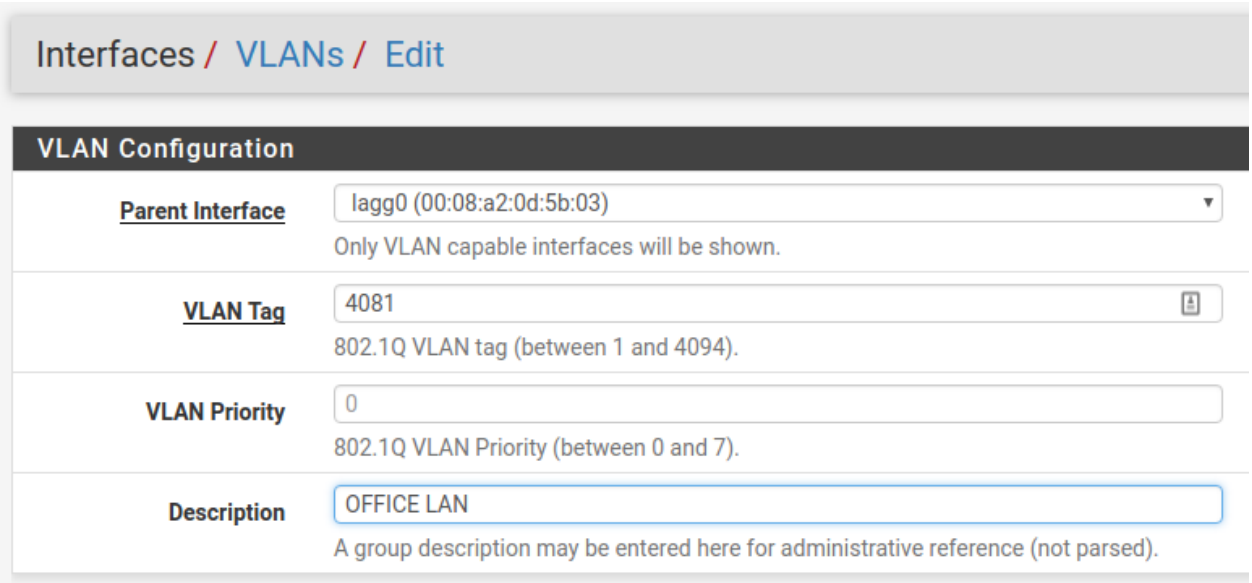

Once the VLAN has been created, it should look something like this:

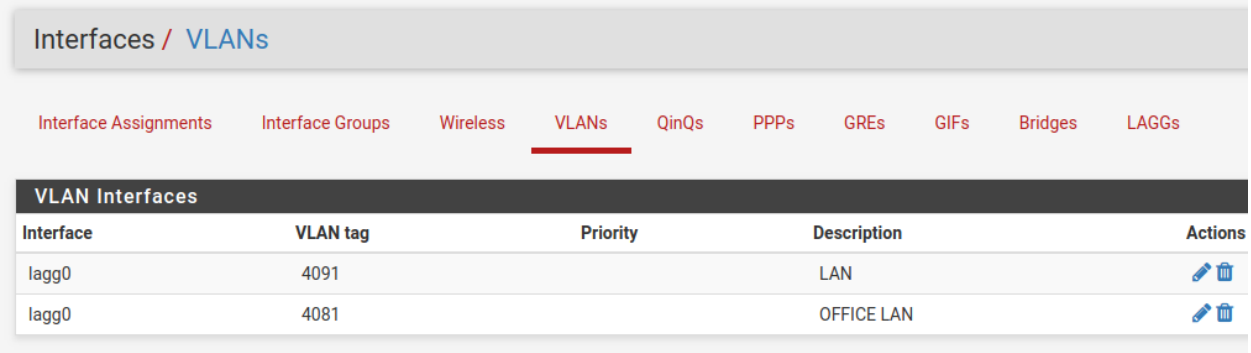

Add, enable, and configure the VLAN interface under Interfaces Assignments:

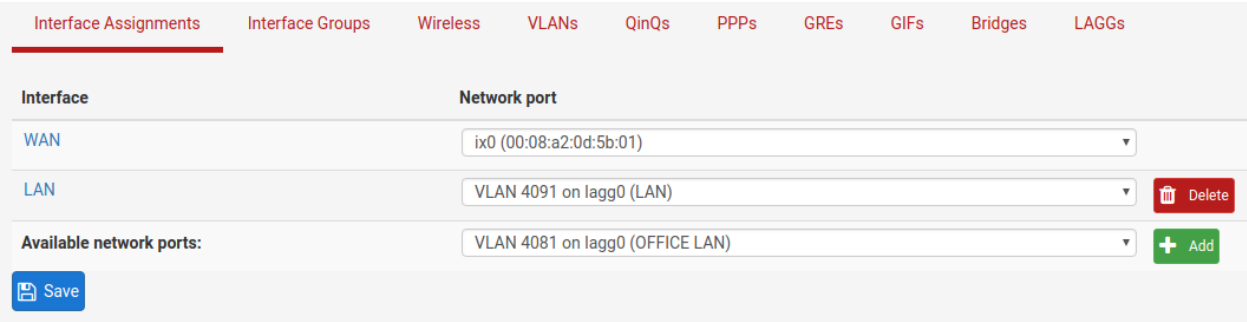

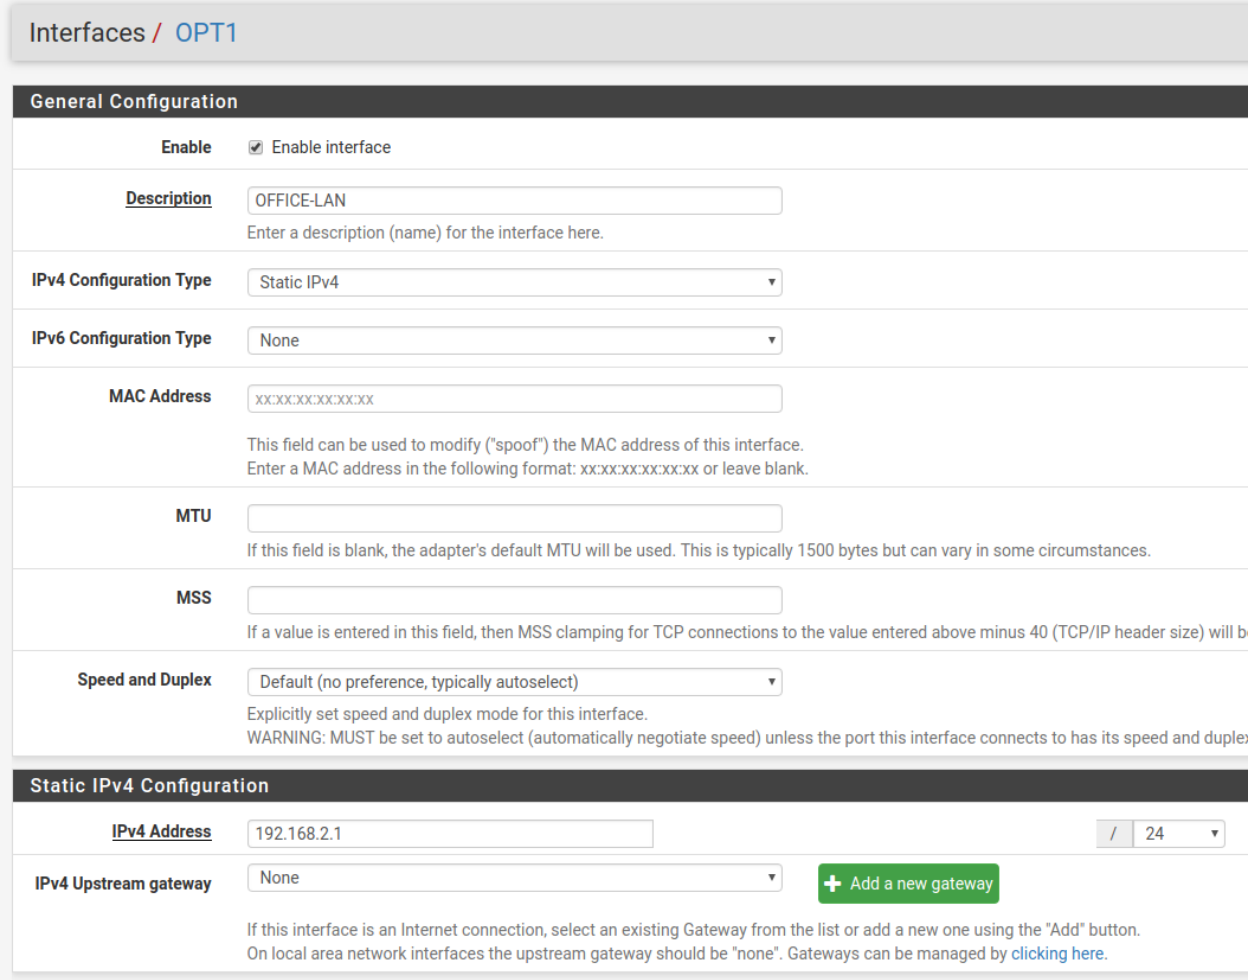

Also create any necessary firewall rules under **Firewall -> Rules**.

Now that pfSense knows of this new VLAN network, configure the switch so that ETH1-4 use the new network. To do this, go to **Interfaces -> Switches -> VLANs** and click the **Add Tag** button. Input the VLAN tag for the new network (same as the VLAN ID configured in the previous steps) and add ETH1-4 and PORT9-10 (uplinks) as members. Be sure **9** and **10** are marked as **tagged**:

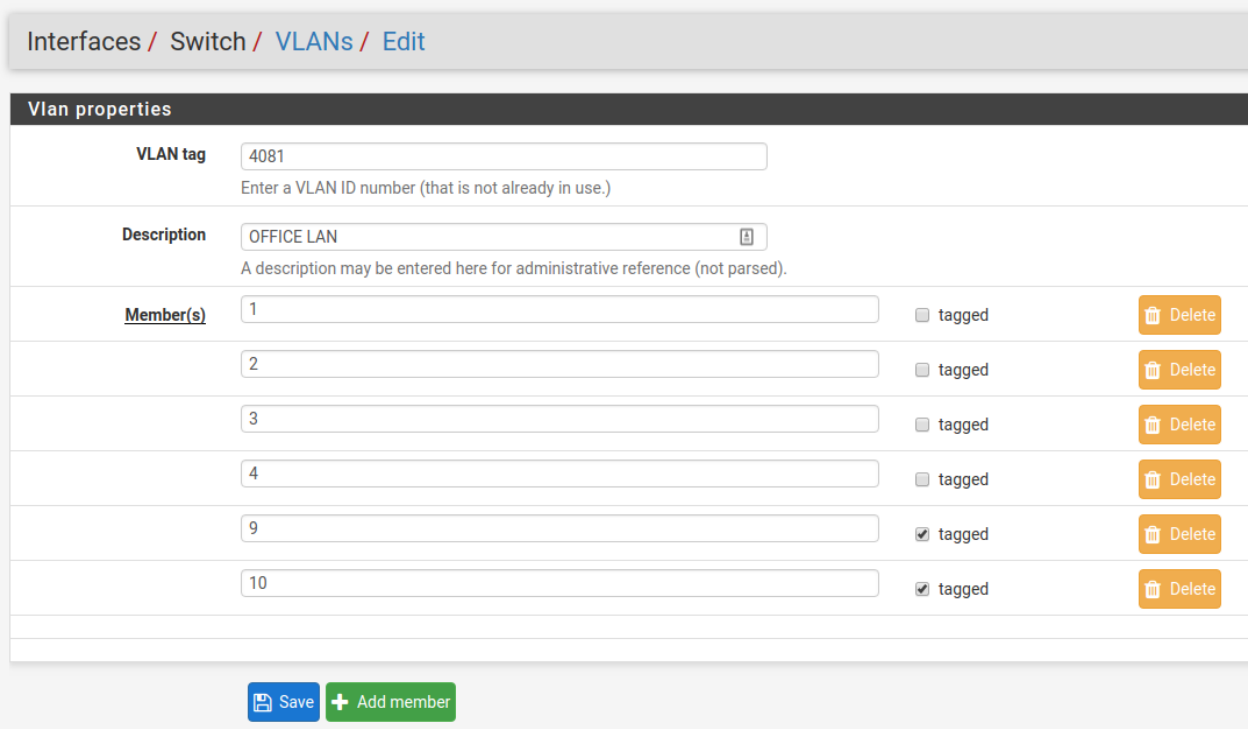

Once this is done, click the **Save** button. The final result should look like this:

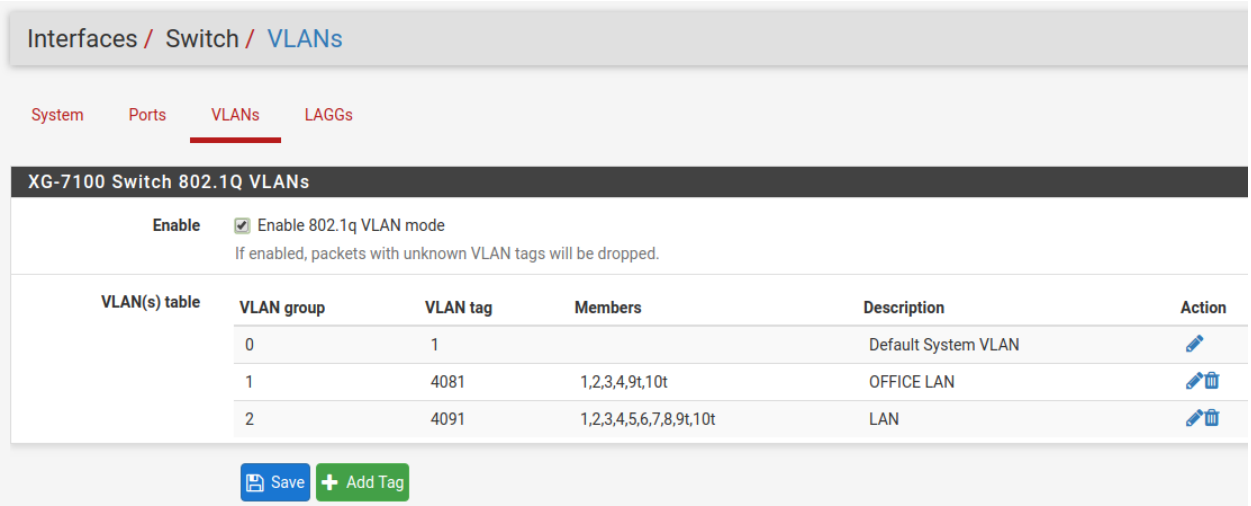

Lastly, update the Port VIDs to use the new 4081 VLAN rather than 4091 on ETH1-4 and click **Save**:

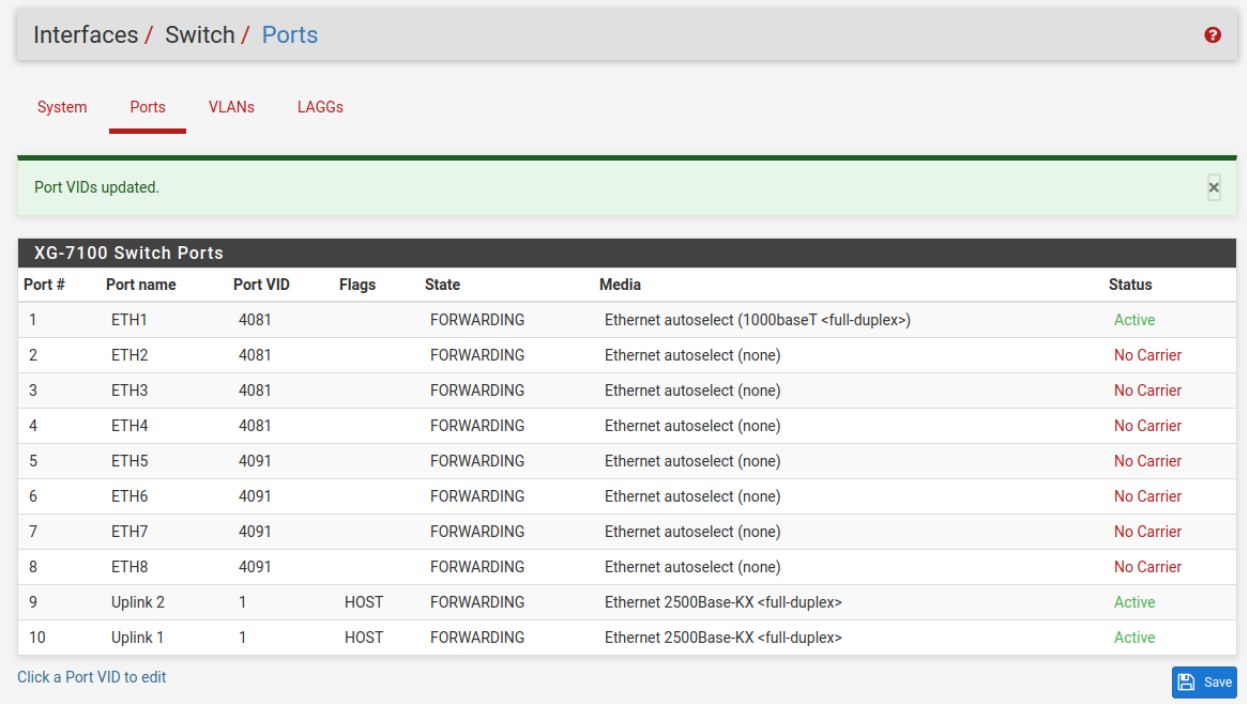

Now ETH1-4 act as a switch for the VLAN 4081 LAN and ETH5-8 act as a switch for the VLAN 4091 LAN.

2.5.3 Trunking VLAN tagged traffic

For expanding on the previous example, let's assume there is a management VLAN of 4000 where devices are already tagged on this VLAN prior to hitting pfSense. Devices on this VLAN may come through on ETH8 but there may also be untagged client traffic.

First, create the management VLAN of 4000 in pfSense using the same steps in the previous example (up to the switch configuration part). Next, add the VLAN to the switch under **Interfaces -> Switches -> VLANs**. ETH8 and PORT9-10 should be added as members and all three will be marked as **tagged**:

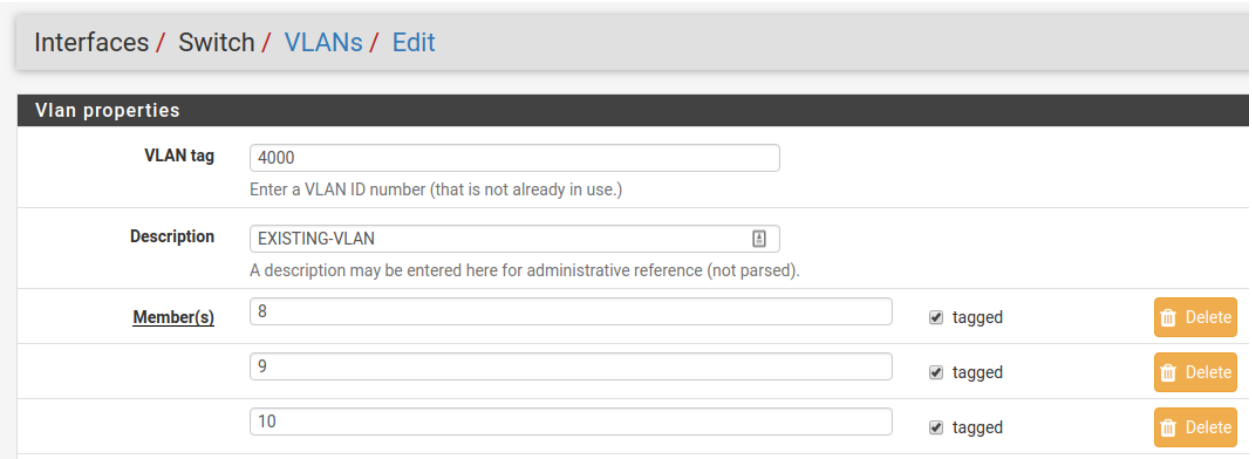

Once it's added, the final result should look like this:

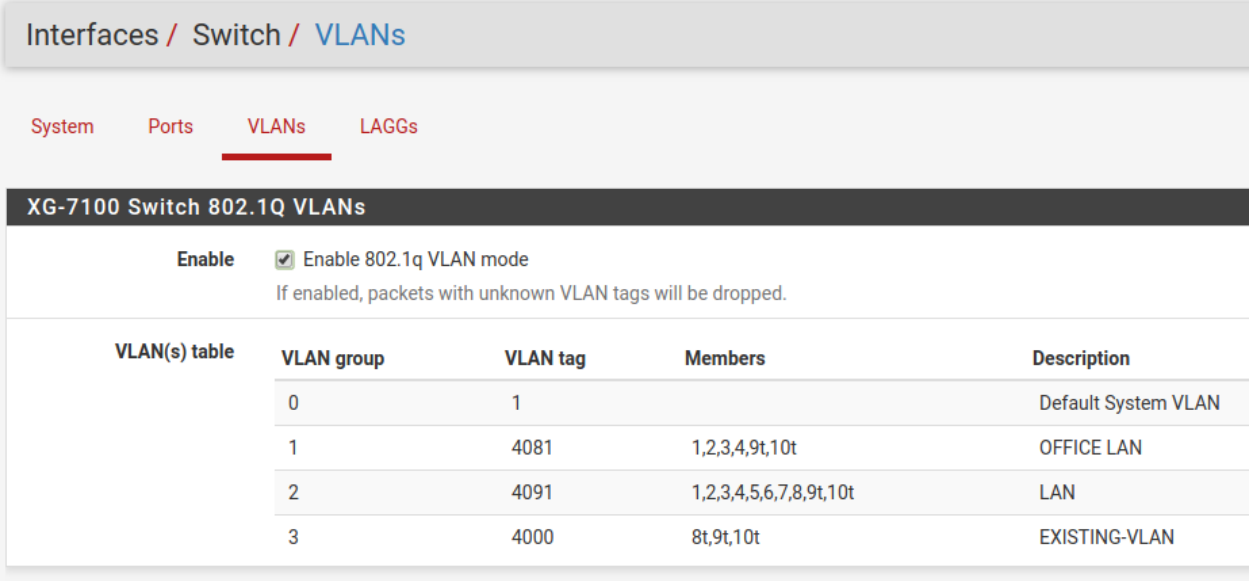

Untagged traffic on ETH8 will be assigned a VLAN ID of 4091. ETH8 and the uplinks will also accept traffic that has already been tagged with a VLAN ID of 4000 as well.

CHAPTER

THREE

GETTING STARTED

Tip: Before configuring the pfSense appliance it is best to activate it by following the instructions at https://www.netgate.com/register/.

The basic firewall configuration begins with connecting the pfSense appliance to the Internet. Neither the modem nor the pfSense appliance should be powered up at this time.

Establishing a connection to the Internet Service Provider (ISP) starts with connecting one end of an ethernet cable to the WAN port (shown in the *I/O Ports* section) of the pfSense appliance.

Warning: The default LAN subnet on the firewall is `192.168.1.0/24`. The same subnet **cannot** be used on both WAN and LAN, so if the subnet on the WAN side of the firewall is also `192.168.1.0/24`, **disconnect the WAN** interface until the LAN interface has been renumbered to a different subnet.

The opposite end of the same ethernet cable should be inserted in to the LAN port of the ISP-supplied modem. The modem provided by the ISP might have multiple LAN ports. If so, they are usually numbered. For the purpose of this installation, please select port 1.

The next step is to connect the LAN port (shown in the *I/O Ports* section) of the pfSense appliance to the computer which will be used to access the firewall console.

Connect one end of the second ethernet cable to the LAN port (shown in the *I/O Ports* section) of the pfSense appliance. Connect the other end to the network connection on the computer. In order to access the web configurator, the PC network interface must be set to use DHCP, or have a static IP set in the `192.168.1.x` subnet with a subnet mask of `255.255.255.0`. Do not use `192.168.1.1`, as this is the address of the firewall, and will cause an IP conflict.

3.1 Initial Setup

The next step is to power up the modem and the firewall. Plug in the power supply to the power port (shown in the *I/O Ports* section).

Once the modem and pfSense appliance are powered up, the next step is to power up the computer.

Once the pfSense appliance is booted, the attached computer should receive a `192.168.1.x` IP address via DHCP from the pfSense appliance.

3.2 Logging Into the Web Interface

Browse to https://192.168.1.1 to access the web interface. In some instances, the browser may respond with a message indicating a problem with website security. Below is a typical example in Google Chrome. If this message or similar message is encountered, it is safe to proceed.

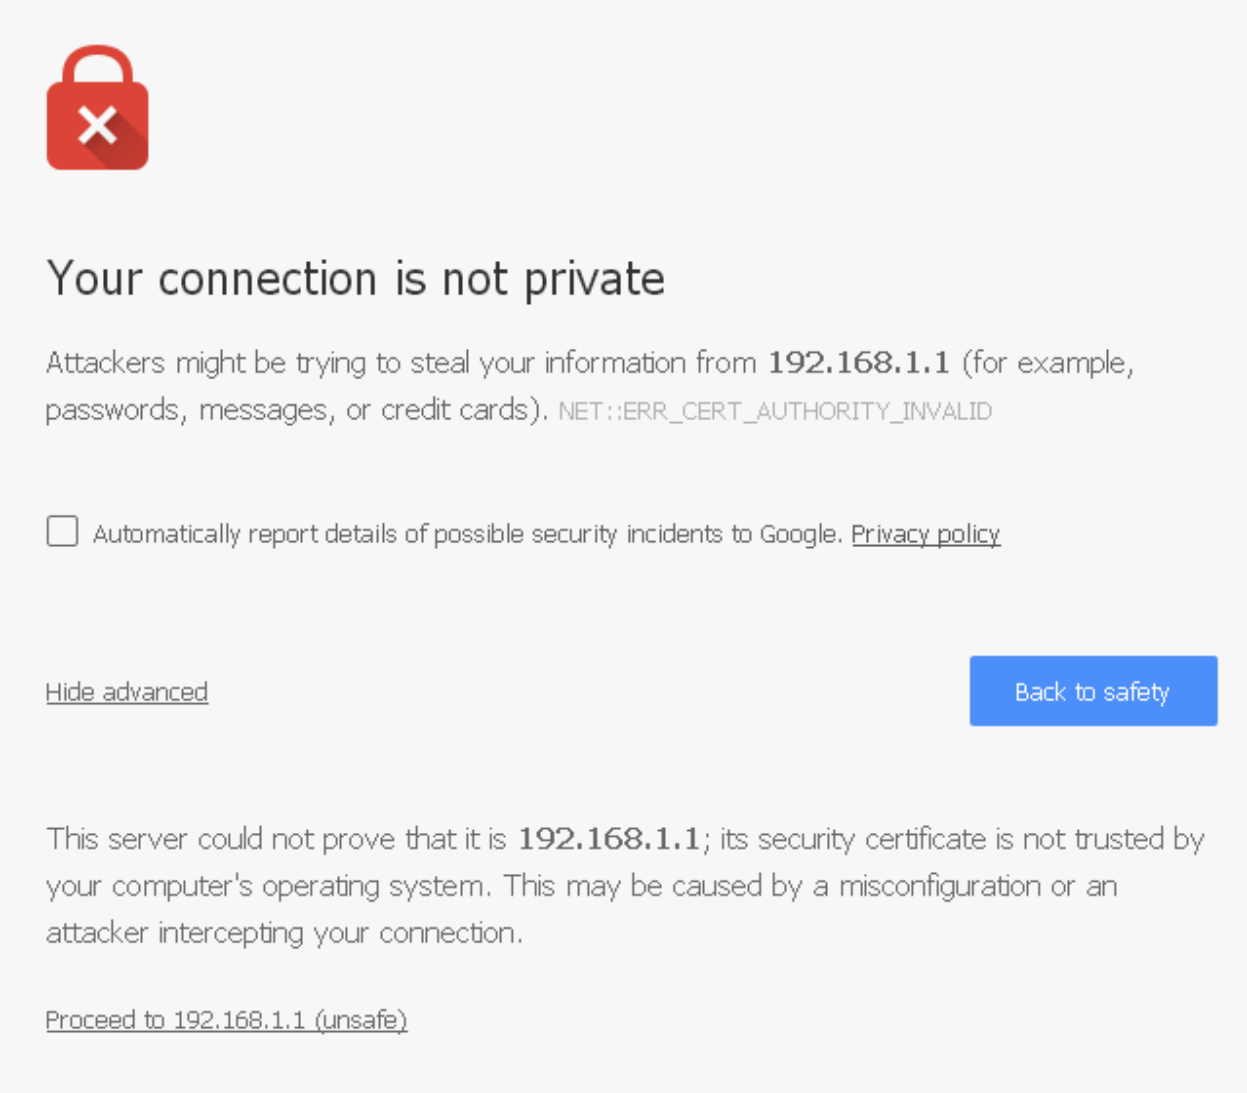

At the login page enter the default pfSense password and username:

Username `admin`

Password `pfsense`

Click **Login** to continue

3.3 Wizard

Upon successful login, the following is displayed.

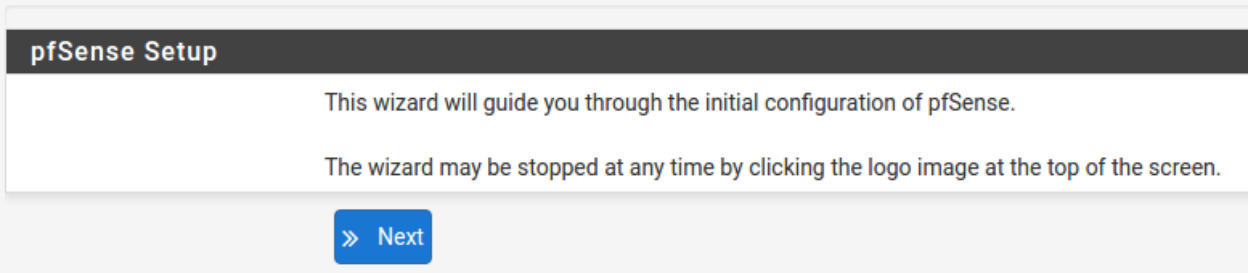

3.4 Configuring Hostname, Domain Name and DNS Servers

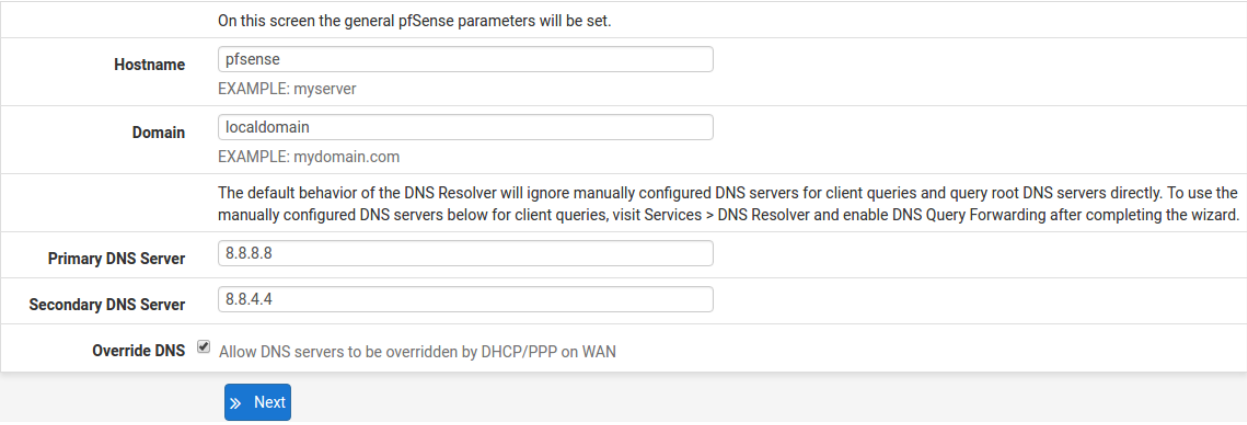

3.5 Hostname

For **Hostname**, any desired name can be entered as it does not affect functionality of the firewall. Assigning a hostname to the firewall will allow the GUI to be accessed by hostname as well as IP address.

For the purposes of this guide, use `pfsense` for the hostname. The default hostname, `pfsense` may be left unchanged.

Once saved in the configuration, the GUI may be accessed by entering http://pfsense as well as http://192.168.1.1

3.6 Domain

If an existing DNS domain is in use within the local network (such as a Microsoft Active Directory domain), use that domain here. This is the domain suffix assigned to DHCP clients, which should match the internal network.

For networks without any internal DNS domains, enter any desired domain name. The default `localdomain` is used for the purposes of this tutorial.

3.7 DNS Servers

The DNS server fields can be left blank if the DNS Resolver is used in non- forwarding mode, which is the default behavior. The settings may also be left blank if the WAN connection is using DHCP, PPTP or PPPoE types of Internet

connections and the ISP automatically assigns DNS server IP addresses. When using a static IP on WAN, DNS server IP addresses must be entered here for name resolution to function if the default DNS Resolver settings are not used.

DNS servers can be specified here even if they differ from the servers assigned by the ISP. Either enter the IP addresses provided by the ISP, or consider using Google public DNS servers (`8.8.8.8`, `8.8.4.4`). Google DNS servers are used for the purpose of this tutorial. Click **Next** after filling in the fields as appropriate.

3.8 Time Server Configuration

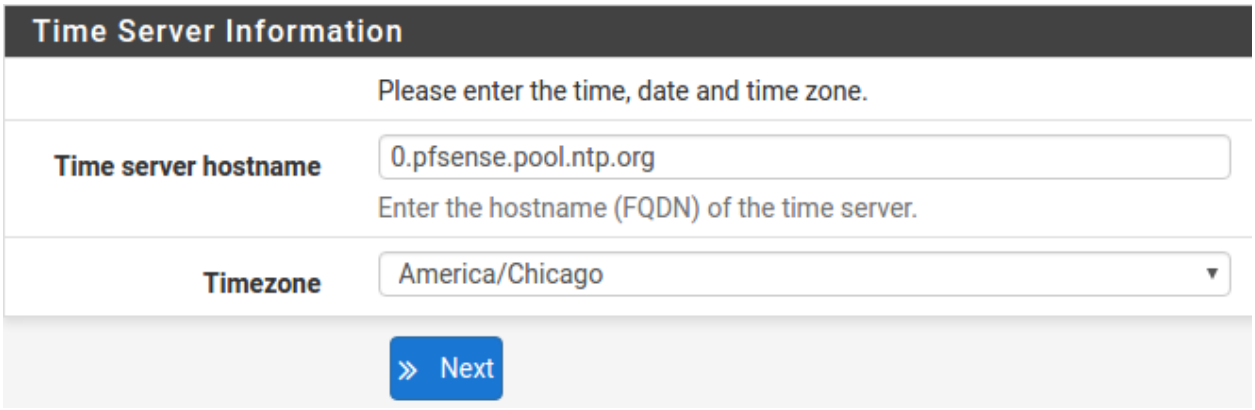

3.9 Time Server Synchronization

Setting time server synchronization is quite simple. We recommend using the default pfSense time server address, which will randomly select an NTP server from a pool.

3.10 Setting Time Zone

Select an appropriate time zone for the location of the firewall. For purposes of this manual, the Timezone setting will be set to `America/Chicago` for US Central time.

3.11 Configuring Wide Area Network (WAN) Type

The WAN interface type is the next to be configured. The IP address assigned to this section becomes the Public IP address that this network will use to communicate with the Internet.

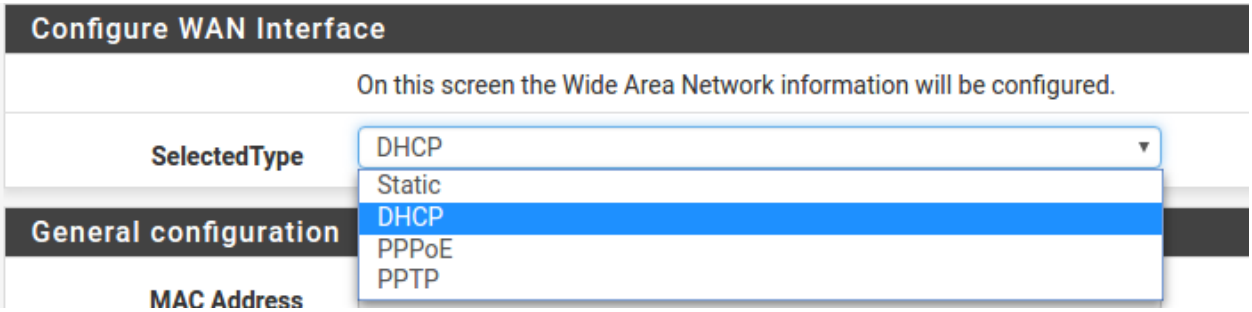

This depicts the four possible WAN interface types. Static, DHCP, PPPoE and PPTP. One must be selected from the drop-down list.

Further information from the ISP is required to proceed when selecting *Static*, *PPPoE* and *PPTP* such as login name and password or as with static addresses, an IP address, subnet mask and gateway address.

DHCP is the most common type of interface for home cable modems. One dynamic IP address is issued from the ISP DHCP server and will become the public IP address of the network behind this firewall. This address will change periodically at the discretion of the ISP. Select *DHCP* as shown and proceed to the next section.

3.12 MAC Address

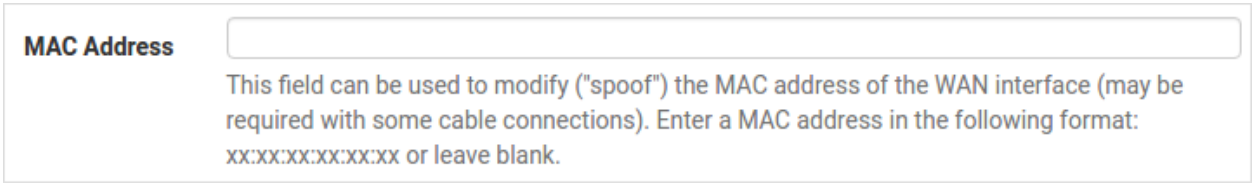

If replacing an existing firewall, the WAN MAC address of the old firewall may be entered here, if it can be determined. This can help avoid issues involved in switching out firewalls, such as ARP caches, ISPs locking to single MAC addresses, etc.

If the MAC address of the old firewall cannot be located, the impact is most likely insignificant. Power cycle the ISP router and modem and the new MAC address will usually be able to get online. For some ISPs, it may be necessary to call them when switching devices, or an activation process may be required.

3.13 Configuring MTU and MSS

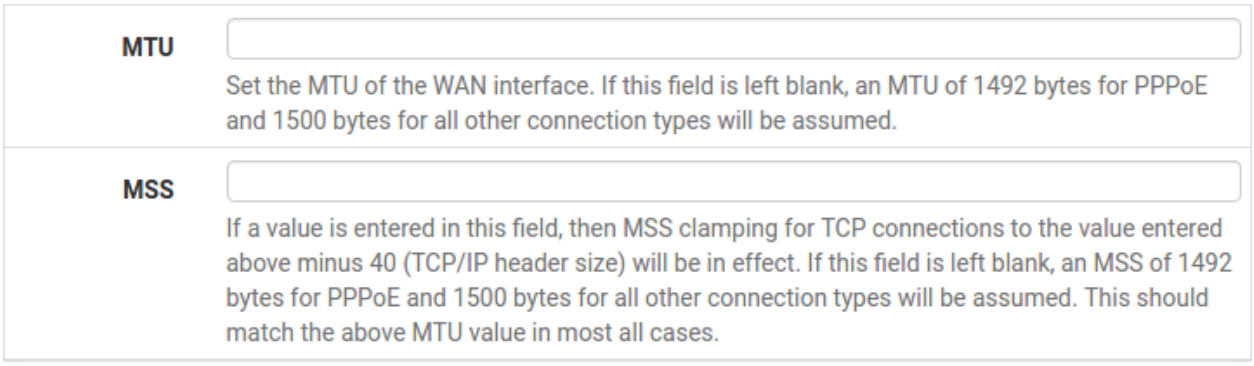

MTU or Maximum Transmission Unit determines the largest protocol data unit that can be passed onwards. A 1500-byte packet is the largest packet size allowed by Ethernet at the network layer and for the most part, the Internet so leaving this field blank allows the system to default to 1500-byte packets. PPPoE is slightly smaller at 1492-bytes. Leave this blank for a basic configuration.

3.14 Configuring DHCP Hostname

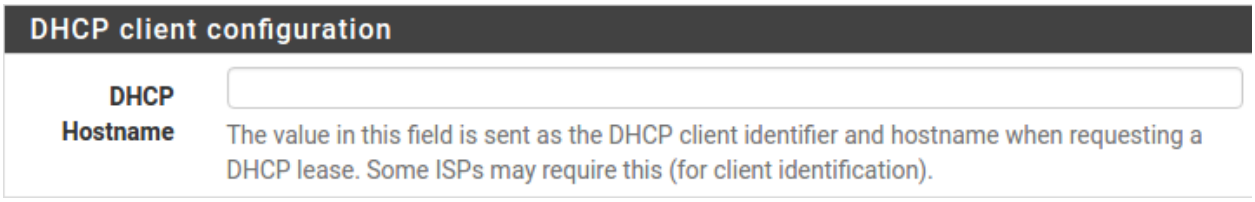

Some ISPs specifically require a **DHCP Hostname** entry. Unless the ISP requires the setting, leave it blank.

3.15 Configuring PPPoE and PPTP Interfaces

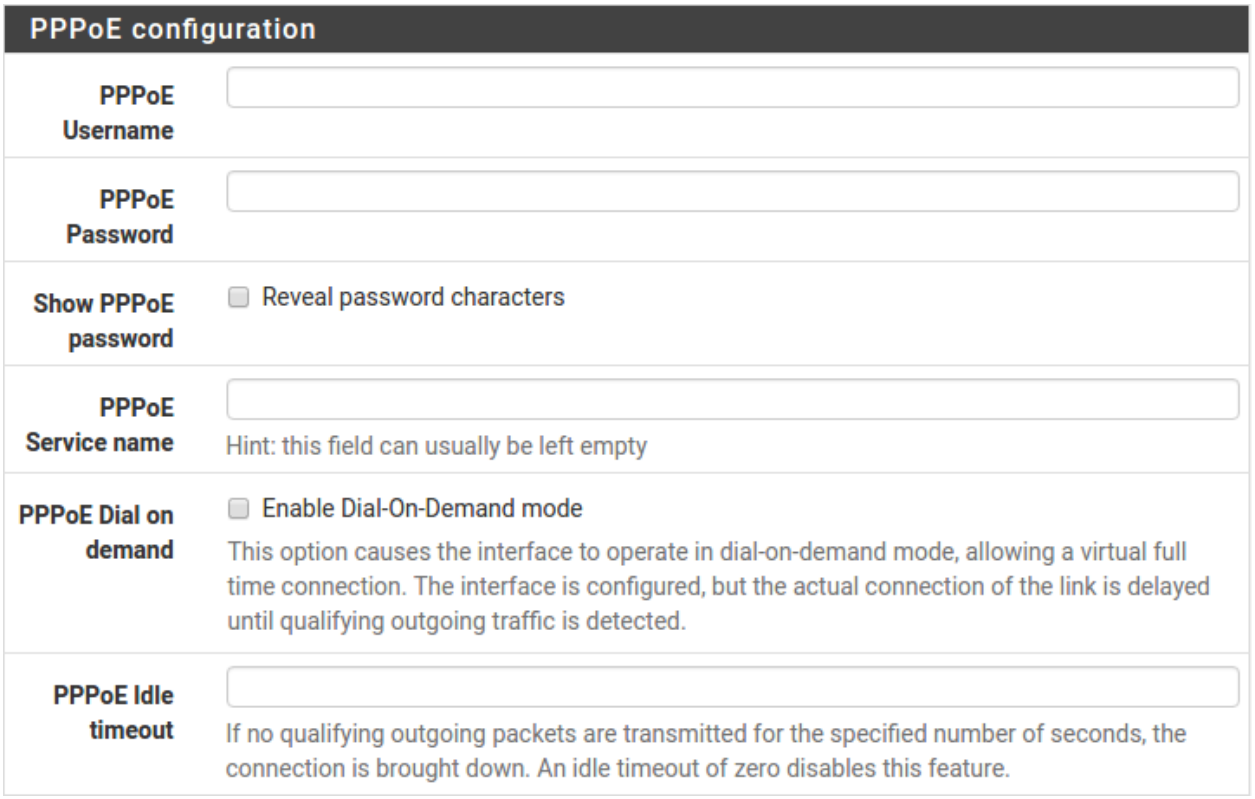

Information added in these sections is assigned by the ISP. Configure these settings as directed by the ISP

3.16 Block Private Networks and Bogons

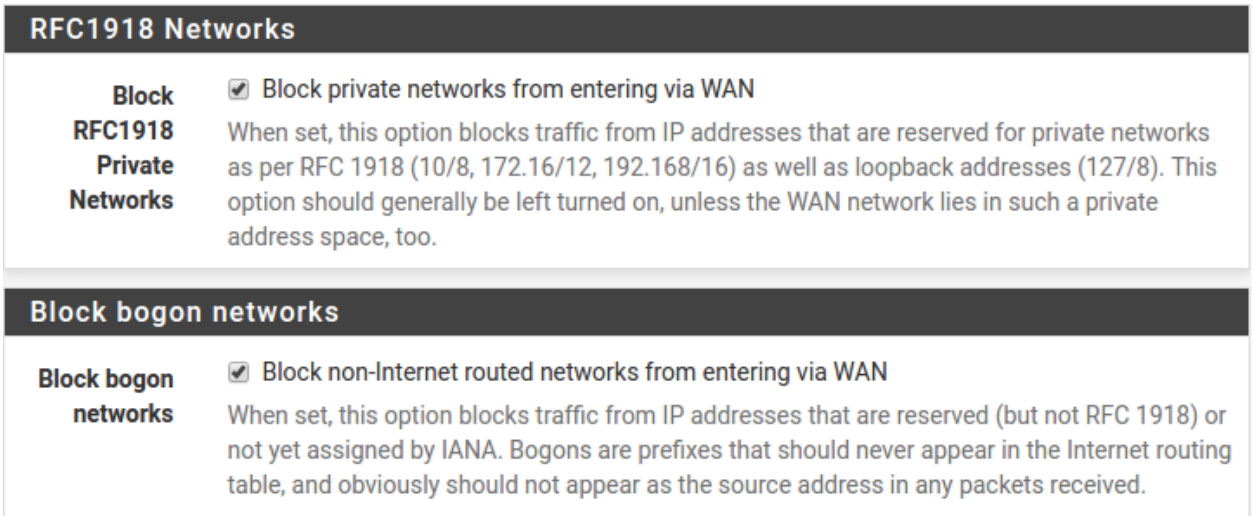

When enabled, all private network traffic originating on the internet is blocked.

Private addresses are reserved for use on internal LANs and blocked from outside traffic so these address ranges may be reused by all private networks.

The following inbound address Ranges are blocked by this firewall rule:

- `10.0.0.1` to `10.255.255.255`
- `172.16.0.1` to `172.31.255.254`
- `192.168.0.1` to `192.168.255.254`
- `127.0.0.0/8`
- `100.64.0.0/10`
- `fc00::/7`

Bogons are public IP addresses that have not yet been allocated, so they may typically also be safely blocked as they should not be in active use.

Check **Block RFC1918 Private Networks** and **Block Bogon Networks**.

Click **Next** to continue.

3.17 Configuring LAN IP Address & Subnet Mask

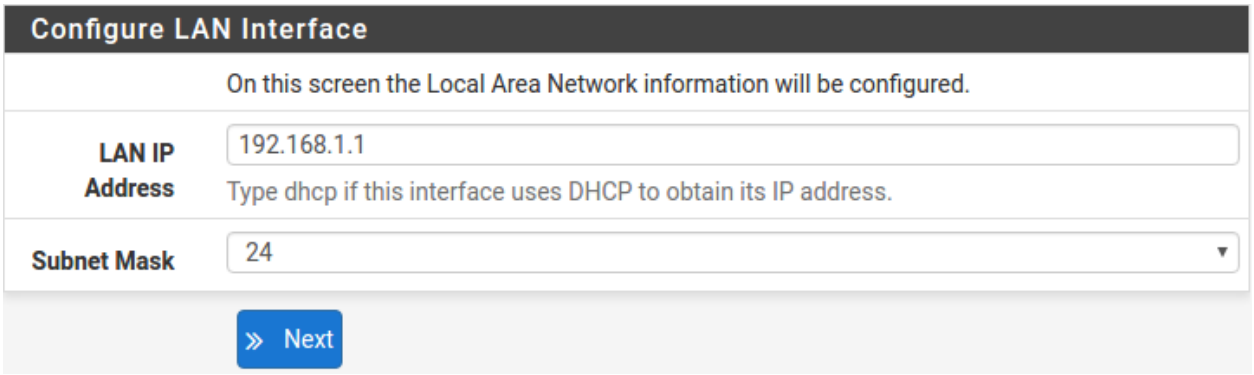

A static IP address of `192.168.1.1` and a subnet mask (CIDR) of `24` was chosen for this installation. If there are no plans to connect this network to any other network via VPN, the `192.168.1.x` default is sufficient.

Click **Next** to continue.

Note: If a Virtual Private Network (VPN) is configured to remote locations, choose a private IP address range more obscure than the very common `192.168.1.0/24`. IP addresses within the `172.16.0.0/12` RFC1918 private address block are the least frequently used. We recommend selecting a block of addresses between `172.16.x.x` and `172.31.x.x` for least likelihood of having VPN connectivity difficulties. An example of a conflict would be If the local LAN is set to `192.168.1.x` and a remote user is connected to a wireless hotspot using `192.168.1.x` (very common), the remote client won't be able to communicate across the VPN to the local network.

3.18 Change Administrator Password

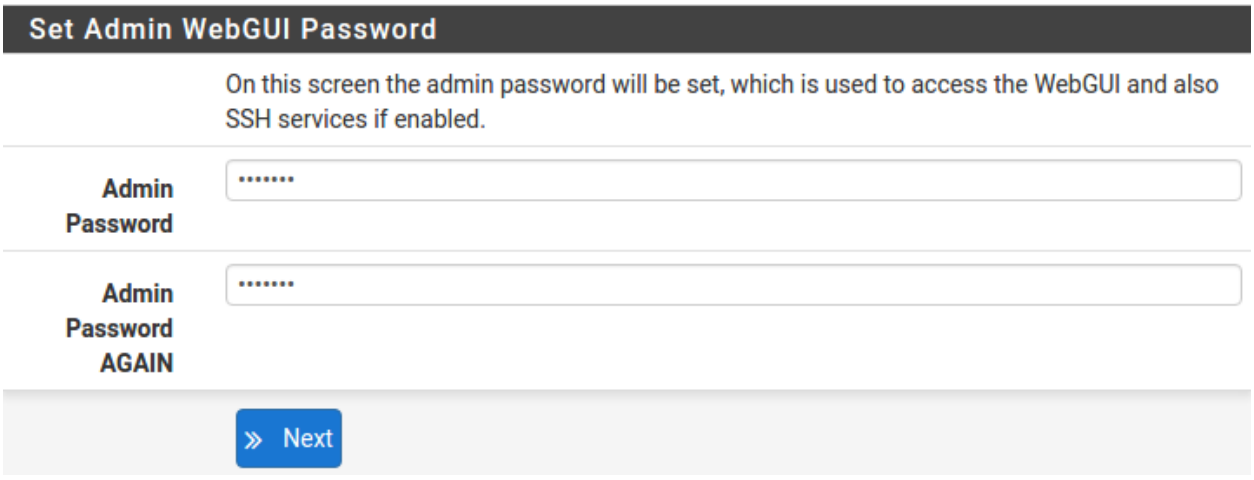

Select a new **Administrator Password** and enter it twice, then click **Next** to continue.

3.19 Save Changes

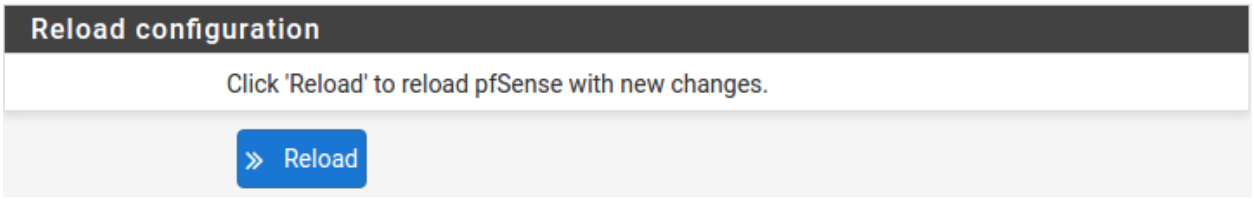

Click **Reload** to save configuration.

3.20 Basic Firewall Configured

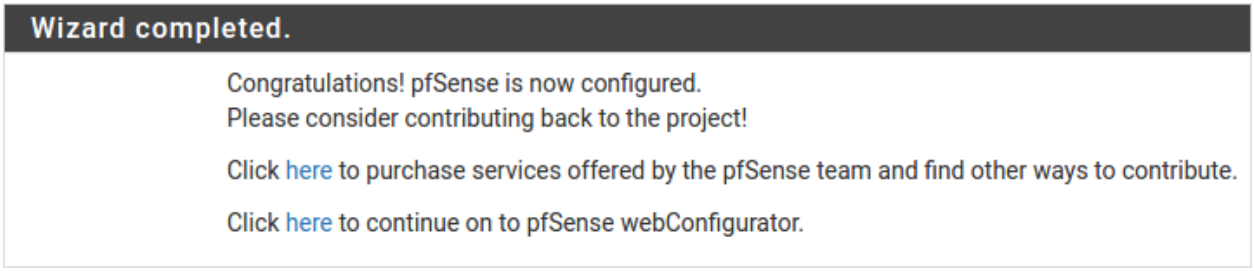

To proceed to the webConfigurator, make the selection as highlighted. The Dashboard display will follow.

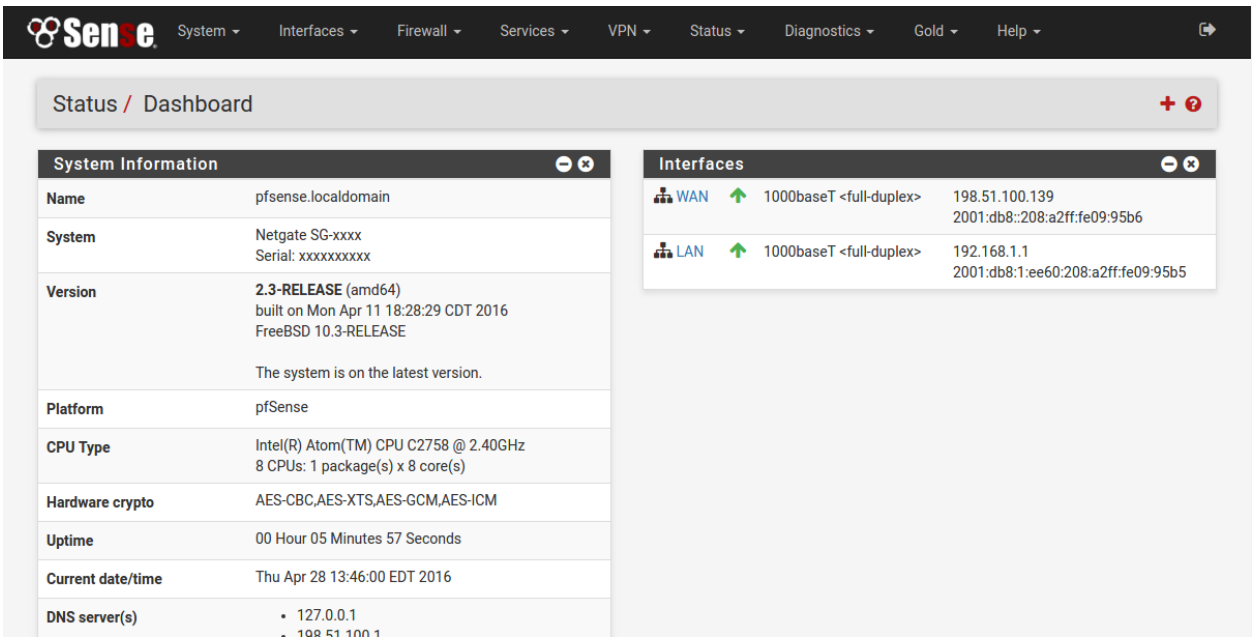

3.21 Backing Up and Restoring

At this point, basic LAN and WAN interface configuration is complete. Before proceeding, backup the firewall configuration. From the menu at the top of the page, browse to **Diagnostics > Backup/Restore**.

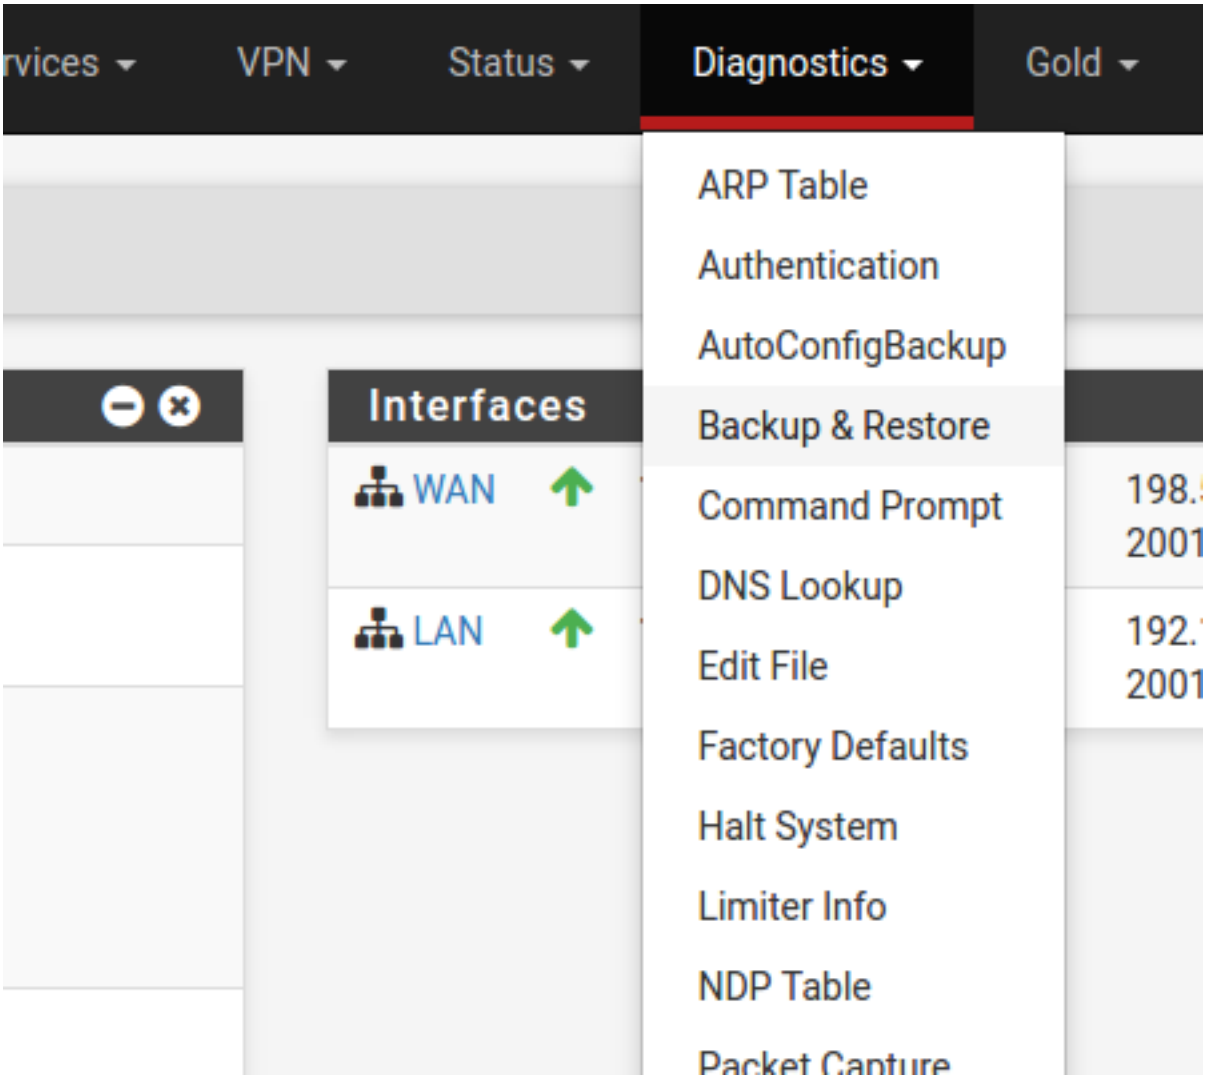

Click **Download Configuration** and save a copy of the firewall configuration.

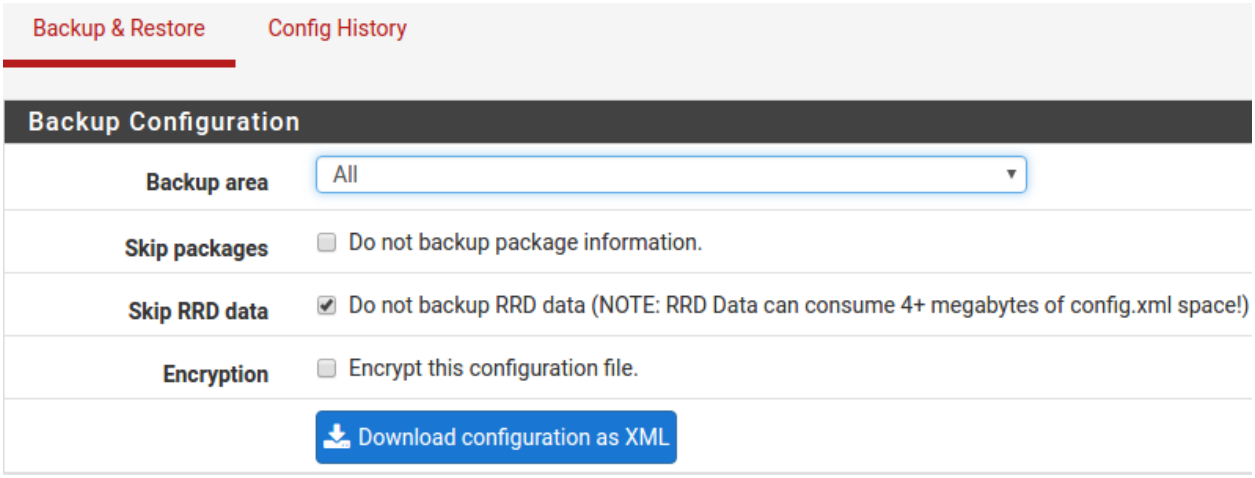

This configuration can be restored from the same screen by choosing the backup file under **Restore configuration**.

3.22 Connecting to the Console

There are times when accessing the console is required. Perhaps GUI console access has been locked out, or the password has been lost or forgotten.

See also:

Connecting to Console Port Connect to the console. Cable is required.

CHAPTER

FOUR

CONNECTING TO CONSOLE PORT

4.1 Simple Configuration

Below are the simple instructions for connecting to the console port with Microsoft Windows. If these steps do not work for you or if you're an operating system other than Windows, then please skip forward to *Advanced Configuration*.

4.1.1 Serial Terminal Emulation Client

A serial terminal emulation program is required to access the pfSense appliance console through the serial interface. Microsoft Windows no longer includes HyperTerminal in Versions 7 and up. PuTTY is free and can be downloaded from:

http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

4.1.2 Configuring Serial Terminal Emulator

PuTTY must be configured to communicate with the pfSense appliance. In order to do so, you must first know what COM Port your computer has assigned to your serial port. Even if you assigned your serial port to COM1 in the BIOS, Windows may remap it to a different COM Port.

To determine this, you must open Windows Device Manager and view the COM port assignment:

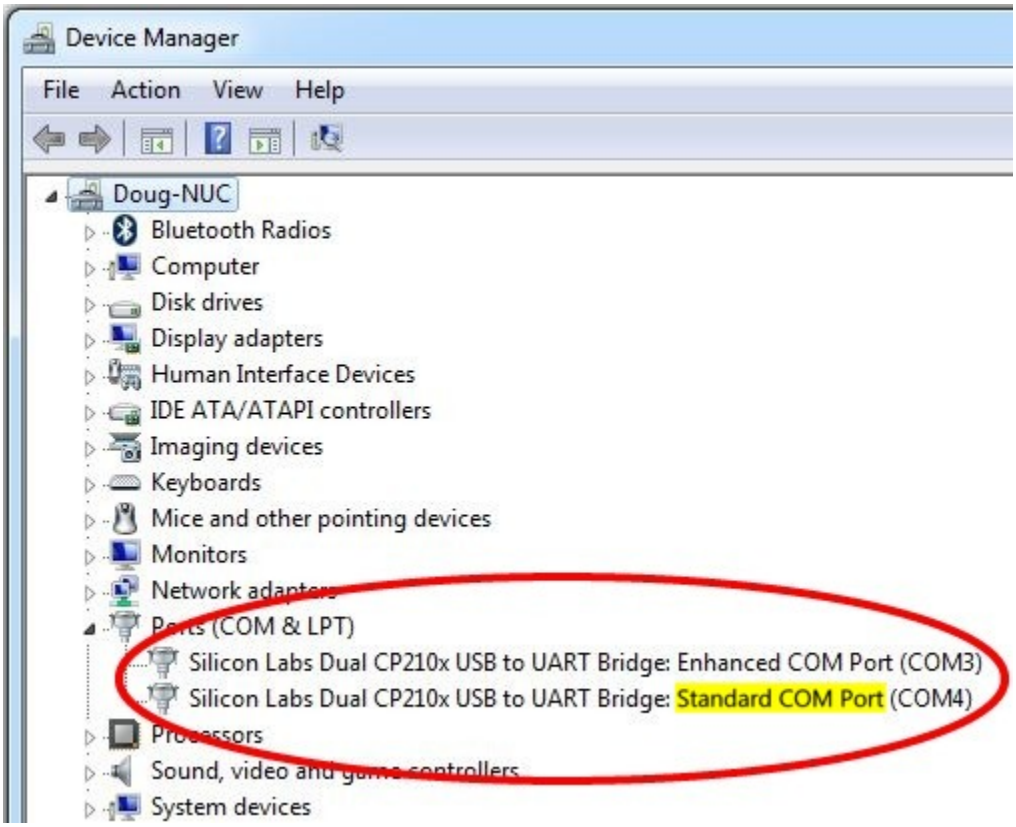

Open PuTTY and locate the **Session** display as shown below. For the **Connection type**, select **Serial**. Set **Serial line** to the COM Port that is displayed in Windows Device Manager, COM4 for this example, and the **Speed** to `115200` bits per second, the speed of the BIOS in this case.

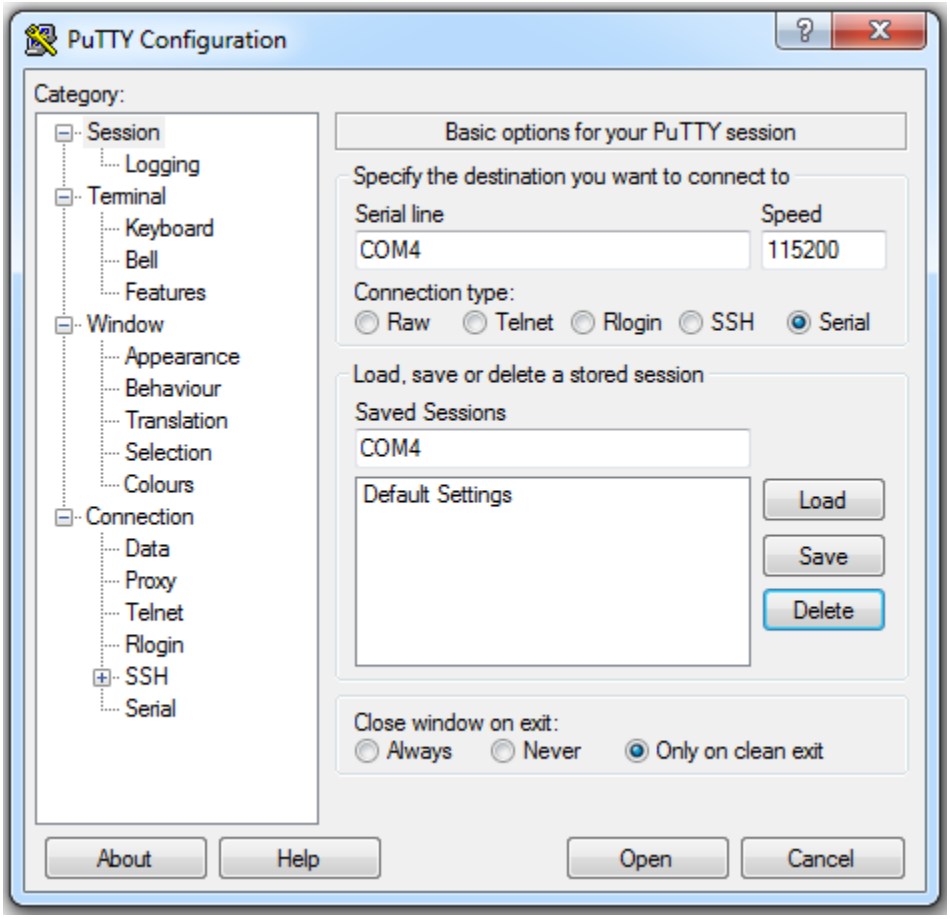

Select **Open** and the console screen will be displayed.

4.2 Advanced Configuration

A Silicon Labs CP210x USB-to-UART bridge is used to provide access to the serial port that acts as a system console. This is exposed via a USB Mini-b (5-pin) port on the front of the case. There are several steps required to access the system console via this port.

4.2.1 Install the Driver

Install an appropriate CP210x USB to UART Bridge VCP (virtual COM port) driver on the workstation used to connect with the system if needed. There are drivers available for Windows, Mac OS X, and Linux available in the Download Software section of the Silicon Labs Website.

Warning: Do not download the **SDK**, only download the driver.

Note: Recent versions of FreeBSD and many Linux distributions include this driver and will not require manual installation.

Loading the Linux Driver

If the device does not appear automatically, the CP210x driver module may need to be loaded manually, especially if the version of Linux being run is not recent. If the driver was provided with the Linux distribution, run `modprobe cp210x` as root or using `sudo`. If it had to be built manually, run `insmod ./cp210x.ko` assuming the module is in the current directory.

4.2.2 Connect a USB Cable

Next, locate an appropriate USB cable. The type of cable required for the serial console has a USB Mini-b (5-pin) connector on one end and a regular USB (Type A) plug on the other end. These cables are commonly used with smaller USB peripherals such as GPS units, cameras, and so on.

Attach the USB cable between a workstation and the system. Gently push the Mini-B plug end into the console port on the system and connect the USB type A plug into an available USB port on the workstation.

Tip: Be certain to gently push in the Mini-B connector on the system side completely. With most cables there will be a tangible "click", "snap", or similar indication when the cable is fully engaged.

4.2.3 Locate the Console Port Device

The appropriate device to attach the terminal program to each platform varies by platform and must be located before attempting to connect to the console.

Windows

To locate the device name on Windows, open **Device Manager** and expand the section for **Ports (COM & LPT)**. Look for an entry with a title such as **Silicon Labs CP210x USB to UART Bridge**. If there is a label in the name that contains "COMX" where X is a decimal digit (e.g. COM1), that value is what would be used as the port in the terminal program.

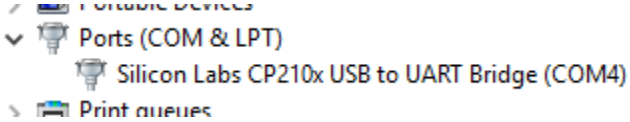

Mac OS X

The device associated with the system console is likely to show up as /dev/cu.SLAB_USBtoUART1.

Linux

The device associated with the system console is likely to show up as /dev/ttyUSB1. Look for messages about the device attaching in the system log files or by running `dmesg`.

Note: If the device does not appear in **/dev/**, see the note above in the driver section about manually loading the Linux driver and then try again.

FreeBSD

The device associated with the system console is likely to show up as /dev/cuaU1. Look for messages about the device attaching in the system log files or by running `dmesg`.

4.2.4 Launch a Terminal Program

Use a terminal program to connect to the system console port. PuTTY is a popular terminal program that is available on various operating systems. Some other choices of terminal programs:

- Linux: screen, PuTTY, minicom, dterm
- Mac OS X: screen, ZTerm, cu
- Windows: PuTTY, SecureCRT, **Do not use Hyperterminal**
- FreeBSD: screen, cu

The settings to use within the terminal program are:

Speed 115200 baud

Data bits 8

Parity none

Stop bits 1

Flow Control Off or XON/OFF. Hardware flow control (RTS/CTS) **must** be disabled.

Client-Specific Examples

PuTTY

Launch PuTTY and configure it for a **Serial** type connection with a speed of **115200** using the port name located previously.

- Windows Example:

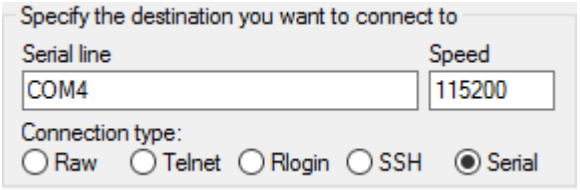

- Linux Example:

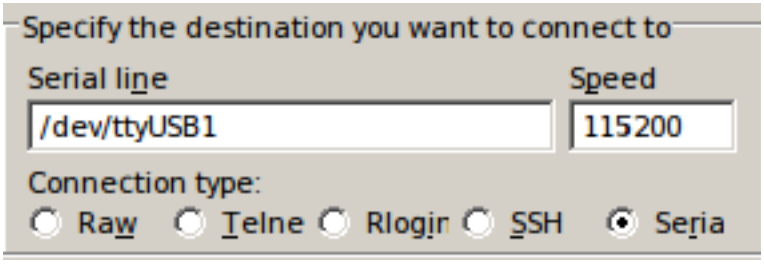

PuTTY generally handles most cases OK but can have issues with line drawing characters on certain platforms.

These settings seem to work best (tested on Windows):

Window Columns x Rows = *80x24*

Window > Appearance Font = *Courier New 10pt* or *Consolas 10pt*

Window > Translation Remote Character Set = *Use font encoding* or *UTF-8*

Window > Translation Handling of line drawing characters = *Use font in both ANSI and OEM modes* or *Use Unicode line drawing code points*

Window > Colours Indicate bolded text by changing = The colour

GNU screen

In many cases *screen* may be invoked simply by using the proper command line:

- Mac OS X

  ```
  sudo screen /dev/cu.SLAB_USBtoUART1 115200
  ```

- Linux

  ```
  sudo screen /dev/ttyUSB1 115200
  ```

- FreeBSD

  ```
  sudo screen /dev/cuaU1 115200
  ```

If portions of the text are unreadable but appear to be properly formatted, the most likely culprit is a character encoding mismatch in the terminal. For example, on OS X this is commonly required

```
sudo screen -U /dev/cu.SLAB_USBtoUART1 115200
```

Adding the `-U` parameter to the *screen* command line arguments forces it to use UTF-8 for character encoding.

4.2.5 Troubleshooting

No Serial Output

If there is no output at all, check the following items:

- Ensure the cable is correctly attached and fully inserted
- Ensure the terminal program is using the correct port
- Ensure the terminal program is configured for the correct speed. The default BIOS speed is 115200, and many other modern operating systems use that speed as well. Some older operating systems or custom configurations may use slower speeds such as 9600 or 38400.
- Ensure the operating system is configured for the proper console (e.g. `ttyS1` in Linux). Consult the various operating install guides on this site for further information.

Garbled Serial Output

If the serial output apears to be garbled, binary, or random characters check the following items:

- Ensure the terminal program is configured for the correct speed. (See "No Serial Output" above)
- Ensure the terminal program is configured for the proper character encoding, such as UTF-8 or Latin-1, depending on the operating system. (See the previous entry under "GNU screen")

Serial Output Stops After the BIOS

If serial output is shown for the BIOS but stops afterward, check the following items:

- Ensure the terminal program is configured for the correct speed for the installed operating system. (See "No Serial Output" above)
- Ensure the installed operating system is configured to activate the serial console.
- Ensure the installed operating system is configured for the proper console (e.g. `ttyS1` in Linux). Consult the various operating install guides on this site for further information.
- If booting from a USB flash drive, ensure that the drive was written correctly and contains a bootable operating system image.

CHAPTER

FIVE

ADDITIONAL RESOURCES

5.1 Professional Services

Support does not cover more complex tasks such as CARP configuration for redundancy on multiple firewalls or circuits, network design, and conversion from other firewalls to pfSense. These items are offered as professional services and can be purchased and scheduled accordingly.

Please see https://www.netgate.com/our-services/professional-services.html for more details

5.2 Netgate Training

Netgate training offers training courses for increasing your knowledge of pfSense products and services. Whether you need to maintain or improve the security skills of your staff or offer highly specialized support and improve your customer satisfaction; Netgate training has got you covered. Check us out at https://www.netgate.com/training/

5.3 Community Support Options

You can find out more information about our active forums, subreddit, IRC, mailing lists and more here: https://www.netgate.com/support/contact-support.html#community-support

CHAPTER

SIX

WARRANTY AND SUPPORT INFORMATION

- One year manufacturer's warranty.
- Please contact Netgate for warranty information or view our Product Lifecycle page.
- All Specifications subject to change without notice

For support information, view our support plans.

CHAPTER

SEVEN

SAFETY AND LEGAL

Contents

7.1 Safety Notices

1. Read, follow, and keep these instructions.
2. Heed all warnings.
3. Only use attachments/accessories specified by the manufacturer

Warning: Do not use this product in location that can be submerged by water.

Warning: Do not use this product during an electrical storm to avoid electrical shock.

7.2 Electrical Safety Information

1. Compliance is required with respect to voltage, frequency, and current requirements indicated on the manufacturer's label. Connection to a different power source than those specified may result in improper operation, damage to the equipment or pose a fire hazard if the limitations are not followed.
2. There are no operator serviceable parts inside this equipment. Service should be provided only by a qualified service technician.
3. This equipment is provided with a detachable power cord which has an integral safety ground wire intended for connection to a grounded safety outlet.
   (a) Do not substitute the power cord with one that is not the provided approved type. If a 3 prong plug is provided, never use an adapter plug to connect to a 2-wire outlet as this will defeat the continuity of the grounding wire.
   (b) The equipment requires the use of the ground wire as a part of the safety certification, modification or misuse can provide a shock hazard that can result in serious injury or death.
   (c) Contact a qualified electrician or the manufacturer if there are questions about the installation prior to connecting the equipment.
   (d) Protective grounding/earthing is provided by Listed AC adapter. Building installation shall provide appropriate short-circuit backup protection.
   (e) Protective bonding must be installed in accordance with local national wiring rules and regulations.

7.3 FCC Compliance

Changes or modifications not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment. This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference, and
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a residential environment.

7.4 Industry Canada

This Class B digital apparatus complies with Canadian ICES-3(B). Cet appareil numérique de la classe A est conforme à la norme NMB-(3)B Canada.

7.5 Australia and New Zealand

This is a AMC Compliance level 2 product. This product is suitable for domestic environments.

7.6 CE Marking

CE marking on this product represents the product is in compliance with all directives that are applicable to it.

7.7 RoHS/WEEE Compliance Statement

7.7.1 English

European Directive 2002/96/EC requires that the equipment bearing this symbol on the product and/or its packaging must not be disposed of with unsorted municipal waste. The symbol indicates that this product should be disposed of separately from regular household waste streams. It is your responsibility to dispose of this and other electric and electronic equipment via designated collection facilities appointed by the government or local authorities. Correct disposal and recycling will help prevent potential negative consequences to the environment and human health. For more detailed information about the disposal of your old equipment, please contact your local authorities, waste disposal service, or the shop where you purchased the product.

7.7.2 Deutsch

Die Europäische Richtlinie 2002/96/EC verlangt, dass technische Ausrüstung, die direkt am Gerät und/oder an der Verpackung mit diesem Symbol versehen ist, nicht zusammen mit unsortiertem Gemeindeabfall entsorgt werden darf. Das Symbol weist darauf hin, dass das Produkt von regulärem Haushaltmüll getrennt entsorgt werden sollte. Es liegt in Ihrer Verantwortung, dieses Gerät und andere elektrische und elektronische Geräte über die dafür zuständigen und von der Regierung oder örtlichen Behörden dazu bestimmten Sammelstellen zu entsorgen. Ordnungsgemäßes Entsorgen und Recyceln trägt dazu bei, potentielle negative Folgen für Umwelt und die menschliche Gesundheit zu vermeiden. Wenn Sie weitere Informationen zur Entsorgung Ihrer Altgeräte benötigen, wenden Sie sich bitte an die örtlichen Behörden oder städtischen Entsorgungsdienste oder an den Händler, bei dem Sie das Produkt erworben haben.

7.7.3 Español

La Directiva 2002/96/CE de la UE exige que los equipos que lleven este símbolo en el propio aparato y/o en su embalaje no deben eliminarse junto con otros residuos urbanos no seleccionados. El símbolo indica que el producto en cuestión debe separarse de los residuos domésticos convencionales con vistas a su eliminación. Es responsabilidad suya desechar este y cualesquiera otros aparatos eléctricos y electrónicos a través de los puntos de recogida que ponen a su disposición el gobierno y las autoridades locales. Al desechar y reciclar correctamente estos aparatos estará contribuyendo a evitar posibles consecuencias negativas para el medio ambiente y la salud de las personas. Si desea obtener información más detallada sobre la eliminación segura de su aparato usado, consulte a las autoridades locales, al servicio de recogida y eliminación de residuos de su zona o pregunte en la tienda donde adquirió el producto.

7.7.4 Français

La directive européenne 2002/96/CE exige que l'équipement sur lequel est apposé ce symbole sur le produit et/ou son emballage ne soit pas jeté avec les autres ordures ménagères. Ce symbole indique que le produit doit être éliminé dans un circuit distinct de celui pour les déchets des ménages. Il est de votre responsabilité de jeter ce matériel ainsi que tout autre matériel électrique ou électronique par les moyens de collecte indiqués par le gouvernement et les pouvoirs publics des collectivités territoriales. L'élimination et le recyclage en bonne et due forme ont pour but de lutter contre l'impact néfaste potentiel de ce type de produits sur l'environnement et la santé publique. Pour plus d'informations sur le mode d'élimination de votre ancien équipement, veuillez prendre contact avec les pouvoirs publics locaux, le service de traitement des déchets, ou l'endroit où vous avez acheté le produit.

7.7.5 Italiano

La direttiva europea 2002/96/EC richiede che le apparecchiature contrassegnate con questo simbolo sul prodotto e/o sull'imballaggio non siano smaltite insieme ai rifiuti urbani non differenziati. Il simbolo indica che questo prodotto non deve essere smaltito insieme ai normali rifiuti domestici. È responsabilità del proprietario smaltire sia questi prodotti sia le altre apparecchiature elettriche ed elettroniche mediante le specifiche strutture di raccolta indicate dal governo o dagli enti pubblici locali. Il corretto smaltimento ed il riciclaggio aiuteranno a prevenire conseguenze potenzialmente negative per l'ambiente e per la salute dell'essere umano. Per ricevere informazioni più dettagliate circa lo smaltimento delle vecchie apparecchiature in Vostro possesso, Vi invitiamo a contattare gli enti pubblici di competenza, il servizio di smaltimento rifiuti o il negozio nel quale avete acquistato il prodotto.

7.8 Declaration of Conformity

7.8.1 Česky[Czech]

NETGATE tímto prohla uje, e tento NETGATE device, je ve shod se základními po adavky a dal ími p íslu n mi ustanoveními sm rnice 1999/5/ES.

7.8.2 Dansk [Danish]

Undertegnede NETGATE erklærer herved, at følgende udstyr NETGATE device, overholder de væsentlige krav og øvrige relevante krav i direktiv 1999/5/EF.

7.8.3 Nederlands [Dutch]

Hierbij verklaart NETGATE dat het toestel NETGATE device, in overeenstemming is met de essentiële eisen en de andere relevante bepalingen van richtlijn 1999/5/EG. Bij deze verklaart NETGATE dat deze NETGATE device, voldoet aan de essentiële eisen en aan de overige relevante bepalingen van Richtlijn 1999/5/EC.

7.8.4 English

Hereby, NETGATE , declares that this NETGATE device, is in compliance with the essential requirements and other relevant provisions of Directive 1999/5/EC.

7.8.5 Eesti [Estonian]

Käesolevaga kinnitab NETGATE seadme NETGATE device, vastavust direktiivi 1999/5/EÜ põhinõuetele ja nimetatud direktiivist tulenevatele teistele asjakohastele sätetele.

7.8.6 Suomi [Finnish]

NETGATE vakuuttaa täten että NETGATE device, tyyppinen laite on direktiivin 1999/5/EY oleellisten vaatimusten ja sitä koskevien direktiivin muiden ehtojen mukainen. Français [French] Par la présente NETGATE déclare que l'appareil Netgate, device est conforme aux exigences essentielles et aux autres dispositions pertinentes de la directive 1999/5/CE.

7.8.7 Deutsch [German]

Hiermit erklärt Netgate, dass sich diese NETGATE device, in Übereinstimmung mit den grundlegenden Anforderungen und den anderen relevanten Vorschriften der Richtlinie 1999/5/EG befindet". (BMWi)

7.8.8 Ελληνικ H [Greek]

ΜΕ ΤΗΝ ΠΑΡΟΥΣΑ NETGATE ΔΗΛΩΝΕΙ ΟΤΙ NETGATE device, ΣΥΜΜΟΡΦΩΝΕΤΑΙ ΠΡΟΣ ΤΙΣ ΟΥΣΙΩΔΕΙΣ ΑΠΑΙΤΗΣΕΙΣ ΚΑΙ ΤΙΣ ΛΟΙΠΕΣ ΣΧΕΤΙΚΕΣ ΔΙΑΤΑΞΕΙΣ ΤΗΣ ΟΔΗΓΙΑΣ 1995/5/ΕΚ.

7.8.9 Magyar [Hungarian]

Alulírott, NETGATE nyilatkozom, hogy a NETGATE device, megfelel a vonatkozó alapvetõ követelményeknek és az 1999/5/EC irányelv egyéb elõírásainak.

7.8.10 Íslenska [Icelandic]

Hér me l sir NETGATE yfir ví a NETGATE device, er í samræmi vi grunnkröfur og a rar kröfur, sem ger ar eru í tilskipun 1999/5/EC.

7.8.11 Italiano [Italian]

Con la presente NETGATE dichiara che questo NETGATE device, è conforme ai requisiti essenziali ed alle altre disposizioni pertinenti stabilite dalla direttiva 1999/5/CE.

7.8.12 Latviski [Latvian]

Ar o NETGATE deklar , ka NETGATE device, atbilst Direkt vas 1999/5/EK b tiskaj m pras b m un citiem ar to saist tajiem noteikumiem.

7.8.13 Lietuviškai [Lithuanian]

NETGATE deklaruoja, kad šis NETGATE įrenginys atitinka esminius reikalavimus ir kitas 1999/5/EB Direktyvos nuostatas.

7.8.14 Malti [Maltese]

Hawnhekk, Netgate, jiddikjara li dan NETGATE device, jikkonforma mal- ti ijiet essenzjali u ma provvedimenti o rajn relevanti li hemm fid-Dirrettiva 1999/5/EC.

7.8.15 Norsk [Norwegian]

NETGATE erklærer herved at utstyret NETGATE device, er i samsvar med de grunnleggende krav og øvrige relevante krav i direktiv 1999/5/EF.

7.8.16 Slovensky [Slovak]

NETGATE t mto vyhlasuje, e NETGATE device, sp a základné po iadavky a v etky príslu né ustanovenia Smernice 1999/5/ES.

7.8.17 Svenska [Swedish]

Härmed intygar NETGATE att denna NETGATE device, står I överensstämmelse med de väsentliga egenskapskrav och övriga relevanta bestämmelser som framgår av direktiv 1999/5/EG.

7.8.18 Español [Spanish]

Por medio de la presente NETGATE declara que el NETGATE device, cumple con los requisitos esenciales y cualesquiera otras disposiciones aplicables o exigibles de la Directiva 1999/5/CE.

7.8.19 Polski [Polish]

Niniejszym, firma NETGATE o wiadcza, e produkt serii NETGATE device, spełnia zasadnicze wymagania i inne istotne postanowienia Dyrektywy 1999/5/EC.

7.8.20 Português [Portuguese]

NETGATE declara que este NETGATE device, está conforme com os requisitos essenciais e outras disposições da Directiva 1999/5/CE.

7.8.21 Română [Romanian]

Prin prezenta, NETGATE declară că acest dispozitiv NETGATE este în conformitate cu cerințele esențiale și alte prevederi relevante ale Directivei 1999/5/CE.

7.9 Disputes

ANY DISPUTE OR CLAIM RELATING IN ANY WAY TO YOUR USE OF ANY PRODUCTS/SERVICES, OR TO ANY PRODUCTS OR SERVICES SOLD OR DISTRIBUTED BY RCL OR ESF WILL BE RESOLVED BY BINDING ARBITRATION IN AUSTIN, TEXAS, RATHER THAN IN COURT. The Federal Arbitration Act and federal arbitration law apply to this agreement.

THERE IS NO JUDGE OR JURY IN ARBITRATION, AND COURT REVIEW OF AN ARBITRATION AWARD IS LIMITED. HOWEVER, AN ARBITRATOR CAN AWARD ON AN INDIVIDUAL BASIS THE SAME DAMAGES AND RELIEF AS A COURT (INCLUDING INJUNCTIVE AND DECLARATORY RELIEF OR STATUTORY DAMAGES), AND MUST FOLLOW THE TERMS OF THESE TERMS AND CONDITIONS OF USE AS A COURT WOULD.

To begin an arbitration proceeding, you must send a letter requesting arbitration and describing your claim to the following:

Rubicon Communications LLC
Attn.: Legal Dept.

4616 West Howard Lane, Suite 900
Austin, Texas 78728
legal@netgate.com

The arbitration will be conducted by the American Arbitration Association (AAA) under its rules. The AAA's rules are available at www.adr.org. Payment of all filing, administration and arbitrator fees will be governed by the AAA's rules.

We each agree that any dispute resolution proceedings will be conducted only on an individual basis and not in a class, consolidated or representative action. We also both agree that you or we may bring suit in court to enjoin infringement or other misuse of intellectual property rights.

7.10 Applicable Law

By using any Products/Services, you agree that the Federal Arbitration Act, applicable federal law, and the laws of the state of Texas, without regard to principles of conflict of laws, will govern these terms and conditions of use and any dispute of any sort that might arise between you and RCL and/or ESF. Any claim or cause of action concerning these terms and conditions or use of the RCL and/or ESF website must be brought within one (1) year after the claim or cause of action arises. Exclusive jurisdiction and venue for any dispute or claim arising out of or relating to the parties' relationship, these terms and conditions, or the RCL and/or ESF website, shall be with the arbitrator and/or courts located in Austin, Texas. The judgment of the arbitrator may be enforced by the courts located in Austin, Texas, or any other court having jurisdiction over you.

7.11 Site Policies, Modification, and Severability

Please review our other policies, such as our pricing policy, posted on our websites. These policies also govern your use of Products/Services. We reserve the right to make changes to our site, policies, service terms, and these terms and conditions of use at any time.

7.12 Miscellaneous

If any provision of these terms and conditions of use, or our terms and conditions of sale, are held to be invalid, void or unenforceable, the invalid, void or unenforceable provision shall be modified to the minimum extent necessary in order to render it valid or enforceable and in keeping with the intent of these terms and conditions. If such modification is not possible, the invalid or unenforceable provision shall be severed, and the remaining terms and conditions shall be enforced as written. Headings are for reference purposes only and in no way define, limit, construe or describe the scope or extent of such section. Our failure to act with respect to a breach by you or others does not waive our right to act with respect to subsequent or similar breaches. These terms and conditions set forth the entire understanding and agreement between us with respect to the subject matter hereof, and supersede any prior oral or written agreement pertaining thereto, except as noted above with respect to any conflict between these terms and conditions and our reseller agreement, if the latter is applicable to you.

7.13 Limited Warranty

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITY

THE PRODUCTS/SERVICES AND ALL INFORMATION, CONTENT, MATERIALS, PRODUCTS (INCLUDING SOFTWARE) AND OTHER SERVICES INCLUDED ON OR OTHERWISE MADE AVAILABLE TO YOU THROUGH THE PRODUCTS/SERVICES ARE PROVIDED BY US ON AN "AS IS" AND "AS AVAILABLE" BASIS, UNLESS OTHERWISE SPECIFIED IN WRITING. WE MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, AS TO THE OPERATION OF THE PRODUCTS/SERVICES, OR THE INFORMATION, CONTENT, MATERIALS, PRODUCTS (INCLUDING SOFTWARE) OR OTHER SERVICES INCLUDED ON OR OTHERWISE MADE AVAILABLE TO YOU THROUGH THE PRODUCTS/SERVICES, UNLESS OTHERWISE SPECIFIED IN WRITING. YOU EXPRESSLY AGREE THAT YOUR USE OF THE PRODUCTS/SERVICES IS AT YOUR SOLE RISK.

TO THE FULL EXTENT PERMISSIBLE BY APPLICABLE LAW, RUBICON COMMUNICATIONS, LLC (RCL) AND ELECTRIC SHEEP FENCING (ESF) DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. RCL AND ESF DO NOT WARRANT THAT THE PRODUCTS/SERVICES, INFORMATION, CONTENT, MATERIALS, PRODUCTS (INCLUDING SOFTWARE) OR OTHER SERVICES INCLUDED ON OR OTHERWISE MADE AVAILABLE TO YOU THROUGH THE PRODUCTS/SERVICES, RCL'S OR ESF'S SERVERS OR ELECTRONIC COMMUNICATIONS SENT FROM RCL OR ESF ARE FREE OF VIRUSES OR OTHER HARMFUL COMPONENTS. RCL AND ESF WILL NOT BE LIABLE FOR ANY DAMAGES OF ANY KIND ARISING FROM THE USE OF ANY PRODUCTS/SERVICES, OR FROM ANY INFORMATION, CONTENT, MATERIALS, PRODUCTS (INCLUDING SOFTWARE) OR OTHER SERVICES INCLUDED ON OR OTHERWISE MADE AVAILABLE TO YOU THROUGH ANY PRODUCTS/SERVICES, INCLUDING, BUT NOT LIMITED TO DIRECT, INDIRECT, INCIDENTAL, PUNITIVE, AND CONSEQUENTIAL DAMAGES, UNLESS OTHERWISE SPECIFIED IN WRITING.

IN NO EVENT WILL RCL'S OR ESF'S LIABILITY TO YOU EXCEED THE PURCHASE PRICE PAID FOR THE PRODUCT OR SERVICE THAT IS THE BASIS OF THE CLAIM.

CERTAIN STATE LAWS DO NOT ALLOW LIMITATIONS ON IMPLIED WARRANTIES OR THE EXCLUSION OR LIMITATION OF CERTAIN DAMAGES. IF THESE LAWS APPLY TO YOU, SOME OR ALL OF THE ABOVE DISCLAIMERS, EXCLUSIONS, OR LIMITATIONS MAY NOT APPLY TO YOU, AND YOU MIGHT HAVE ADDITIONAL RIGHTS.

References

- *High Availability*
- *Reinstalling pfSense*
- *BIOS Flash Procedure*
- pfSense Documentation

CHAPTER

EIGHT

HIGH AVAILABILITY

This document covers configuration of a High Availability cluster using the following features:

- CARP for IP address redundancy
- XMLRPC for configuration synchronization
- pfsync for state table synchronization

With this configuration, two units act as an "active/passive" cluster with the primary node working as the master unit and the secondary node in a backup role, taking over as needed if the primary node fails.

8.1 High Availability Prerequisites

Before a redundant configuration can be achieved, a few prerequisites must be met.

8.1.1 Assumptions

This guide assumes that:

- Only two cluster nodes are used.
- Both cluster nodes are the same model with identical hardware specs.
- Both units have a factory default configuration and there are no existing settings on these units.

Warning: Do not connect the LAN port of both units into the same LAN switch until some basic settings have been applied to each node, which will be done by the end of this section. Otherwise there will be an IP address conflict and communication with each node individually will not be possible until the conflict is resolved.

8.1.2 Determine the Synchronization Interface

One interface on each node will be dedicated for synchronization tasks. This is typically referred to as the "Sync" interface, and it is used for configuration synchronization and pfsync state synchronization. Any available interface may be used. It isn't necessary for it to be a high speed port, but it is necessary to choose the same port on both nodes.

Note: Some call this the "CARP" interface but that is incorrect and very misleading. CARP heartbeats happen on each interface with a CARP VIP; CARP traffic and failover actions do not utilize the Sync interface.

8.1.3 Interface Assignments

Interfaces must be assigned in the same order on all nodes exactly. If the interfaces are not aligned, configuration synchronization and other tasks will not behave correctly. The default configuration has all interfaces assigned by default, as seen in the IO Ports section of the unit's product manual, which makes a good starting point for this guide. If any adjustments have been made to the interface assignments, they must be replicated identically on both nodes.

8.1.4 IP Address Requirements

A High Availability cluster needs three IP addresses in each subnet along with a separate unused subnet for the Sync interface. For WANs, this means that a /29 subnet or larger is required for an optimal configuration. One IP address is used by each node, plus a shared CARP VIP address for failover. The synchronization interface only requires one IP address per node.

The IP addresses used in this guide are shown in the following tables, substitute the real IP addresses as needed.

Table 1: WAN IP Address Assignments

IP Address	Usage
198.51.100.200/24	CARP shared IP address
198.51.100.201/24	Primary node WAN IP address
198.51.100.202/24	Secondary node WAN IP address

Table 2: LAN IP Address Assignments

IP Address	Usage
192.168.1.1/24	CARP shared IP address
192.168.1.2/24	Primary node LAN IP address
192.168.1.3/24	Secondary node LAN IP address

Table 3: Sync IP Address Assignments

IP Address	Usage
172.16.1.2/24	Primary node Sync IP address
172.16.1.3/24	Secondary node Sync IP address

Single address CARP

It is technically possible to configure an interface with a CARP VIP as the only IP address in a given subnet, but it is not generally recommended. When used on a WAN, this type of configuration will only allow communication from the primary node to the WAN, which greatly complicates tasks such as updates, package installations, gateway monitoring, or anything that requires external connectivity from the secondary node. It can be a better fit for an internal interface, however internal interfaces do not typically suffer from the same IP address limitations as a WAN, so it is still preferable to configure IP addresses on all nodes. Such a configuration is not covered in this guide.

8.1.5 Determine CARP VHID Availability

CARP can interfere with VRRP, HSRP, or other systems using CARP if conflicting identifiers are used. In order to ensure that a segment is clear of conflicting traffic, perform a packet capture on each interface looking for CARP/VRRP traffic. A given VHID must be unique on each layer 2, so each interface must be checked separately. The same VHID may be used on different segments so long as they are separate broadcast domains.

If any CARP or VRRP traffic is shown, note the VHID/VRID and avoid using that identifier when configuring the CARP VIP VHIDs later.

This guide assumes there is no other potentially conflicting traffic present.

8.1.6 Setup Requirements

Using the Setup Wizard, or manually afterward, configure each firewall with a unique hostname and non-conflicting static IP addresses.

For example, one node could be "firewall-a.example.com" and the other "firewall- b.example.com", or a more personalized pair of names. Avoid naming the nodes "master" and "backup" since those are states and not roles, instead they could be named "primary" and "secondary".

For IP addresses, the factory default LAN address is 192.168.1.1. In a High Availability environment, that address would be a CARP VIP instead. Using that subnet, move each node to its own address there, such as 192.168.1.2 for the primary and 192.168.1.3 for the secondary. This layout is shown in *LAN IP Address Assignments*

Once each node has a unique LAN IP address, then both nodes may be plugged into the same LAN switch.

Both nodes must have the GUI running on the same port and protocol. This guide assumes both use HTTPS on port 443.

The admin account cannot be disabled and both nodes must have the same admin account password.

Both nodes must have static IP addresses in the same subnet and have a proper gateway configured on the WAN interface settings.

Both nodes must have DNS configured properly under **System > General Setup**.

Visit **Services > DNS Resolver**. Review the settings and even if nothing has been changed, click **Save** once to ensure the default values are respected.

8.1.7 Switch / Layer 2 Configuration

CARP Concerns

CARP heartbeats utilize multicast and may require special handling on the switches involved with the cluster. Some switches filter, rate limit, or otherwise interfere with multicast in ways that can cause CARP to fail. Also, some switches employ port security methods which may not work properly with CARP.

At a minimum, the switch must:

- Allow Multicast traffic to be sent and received without interference on ports using CARP VIPs.
- Allow traffic to be sent and received using multiple MAC addresses.
- Allow the CARP VIP MAC address to move between ports.

Nearly all problems with CARP failing to properly reflect the expected status are failures of the switch or other layer 2 issues, so be sure the switches are properly configured before continuing.

Port Configuration

Each node must be connected to a common, but separate, layer 2 on each interface. This means that WAN, LAN, and other interfaces must be connected to separate switches or VLANS with each node being connected to the same segments on each.

For example, the WAN ports of each node must connect to the same WAN switch, which then connects to the WAN CPE/Modem/Upstream link. The LAN ports would all connect to the same LAN switch, and so on. The Sync interface may be connected directly between the two nodes without a switch. See *Example High Availability Cluster* for an example connection layout.

8.2 Configuring a HA Cluster

Note: The WAN and LAN should be configured to static addresses prior to configuring a HA Cluster. Please see *High Availability Prerequisites* for IP address details.

This is the heart of the process, making the changes that will link the systems and allow them to function together.

8.2.1 Setup Sync Interface

Before proceeding, the Sync interfaces on the cluster nodes must be configured. *Sync IP Address Assignments* lists the addresses to use for the Sync interfaces on each node.

1. Navigate to **Interfaces** and choose the interface to use on the SYNC port
2. Check **Enable Interface**
3. Enter `SYNC` for the **Description**
4. Set **IPv4 Configuration Type** to *Static IPv4*
5. Set **IPv4 address** to `172.16.1.2` when configuring the primary node, or `172.16.1.3` when configuring the secondary node
6. Select *24* for the subnet mask in the CIDR drop-down next to **IPv4 address**
7. Do not check **Block private networks** or **Block bogon networks**
8. Click **Save**
9. Click **Apply Changes**

Once that procedure has been completed on the primary node, perform it again on the secondary node with the appropriate **IPv4 address** value. Remember they must be the same on both nodes.

After configuring the sync interface, the interface assignments should have one labeled SYNC.

Add Firewall Rules for Synchronization

To complete the Sync interface configuration, firewall rules must be added to both nodes to allow synchronization.

At a minimum, the firewall rules must pass the configuration synchronization traffic (by default, HTTPS on port 443) and pfsync traffic. In most cases, a simple "allow all" style rule is used. For this guide, both will be shown and it will also serve as an indicator that synchronization is working.

On the primary node:

Set up a rule to allow configuration synchronization:

1. Navigate to **Firewall > Rules** on the **SYNC** tab
2. Click  at the top of the list to create a new rule

3. Set **Action** to *Pass*
4. Set **Source** to *SYNC Net*
5. Set **Destination** to *SYNC Address*
6. Set **Destination port range** to `443` or choose *HTTPS (443)* from the drop-down selector
7. Set **Description** to `Allow configuration synchronization`
8. Click **Save**

Set up a rule to allow state synchronization:

1. Click Add at the top of the list to create another new rule
2. Set **Action** to *Pass*
3. Set **Protocol** to *pfsync*
4. Set **Source** to *SYNC Net*
5. Set **Destination** to *any*
6. Set **Description** to `Allow state synchronization`
7. Click **Save**

Set up a rule to allow ICMP echo (ping) for Diagnostics:

1. Click Add at the top of the list to create another new rule
2. Set **Action** to *Pass*
3. Set **Protocol** to *ICMP*
4. Set **Source** to *SYNC Net*
5. Set **Destination** to *SYNC Net*
6. Set **Description** to `Allow ICMP echo (ping) for Diagnostics`
7. Click **Save**
8. Click **Apply Changes**

When complete, the rules will look like the following, which also includes a rule to allow ICMP echo (ping) for diagnostic purposes.

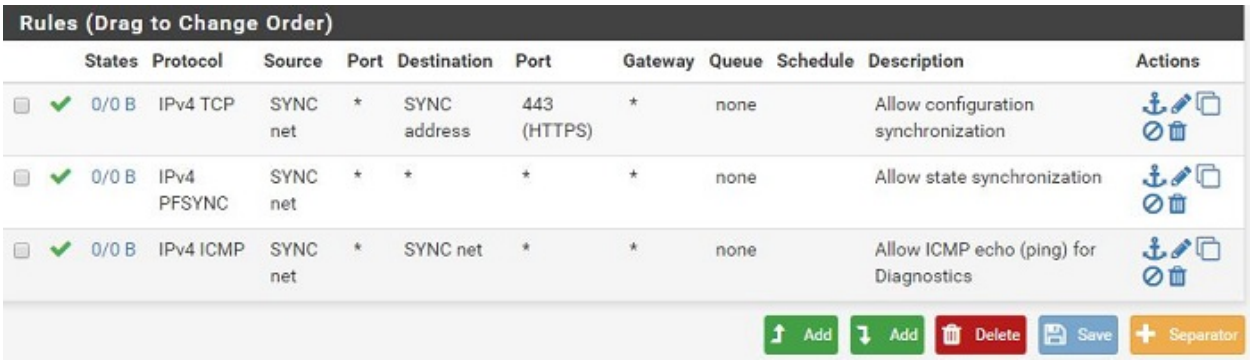

Fig. 1: Example Sync Interface Firewall Rules

On the secondary node:

1. Navigate to **Firewall > Rules** on the **SYNC** tab
2. Click  at the top of the list to create a new rule
3. Set **Action** to *Pass*
4. Set **Protocol** to *any*
5. Set **Source** to *SYNC Net*
6. Set **Destination** to *any*
7. Set **Description** to `Temp rule for sync`
8. Click **Save**
9. Click **Apply Changes**

Note: The rule on the secondary is different, but that is intended at this point. Once the first configuration synchronization has taken place, the temporary rule on the secondary will be replaced by the rules from the primary. Seeing that the rules on the Sync interface changed is a good indicator that it worked!

8.2.2 Configure pfsync

State synchronization using pfsync must be configured on both the primary and secondary nodes to function.

First on the primary node and then on the secondary, perform the following:

1. Navigate to **System > High Avail. Sync**
2. Check **Synchronize States**
3. Set **Synchronize Interface** to *SYNC*
4. Set **pfsync Synchronize Peer IP** to the other node. Set this to `172.16.1.3` when configuring the primary node, or `172.16.1.2` when configuring the secondary node
5. Click **Save**

8.2.3 Configure XMLRPC

Warning: Configuration synchronization **must only be configured on the primary node**. Never activate options in this section on the secondary node of a two-member cluster.

On the primary node only, perform the following:

1. Navigate to **System > High Avail. Sync**
2. Set **Synchronize Config to IP** to the secondary node's Sync interface IP address, `172.16.1.3`
3. Set **Remote System Username** to `admin`.

   Note: This must always be `admin`. No other user will work!

   ---

4. Set **Remote System Password** to the admin user account password and be sure to confirm the password.

5. Check the boxes for each area to synchronize to the secondary node. For this guide, as with most configurations, all boxes are checked.
6. Click **Save**

As a quick confirmation that the synchronization worked, on the secondary node navigate to **Firewall > Rules** on the **SYNC** tab. The rules entered on the primary are now there, and the temporary rule is gone.

The two nodes are now linked for configuration synchronization! Changes made to the primary node in supported areas will be synchronized to the secondary whenever a change is made.

Warning: Do not make changes to the secondary in areas set to be synchronized! These changes will be overwritten the next time the primary node performs a synchronization.

8.2.4 Add CARP VIPs

Now that the configuration synchronization is complete, the CARP Virtual IP addresses need only be added to the primary node and they will be automatically copied to the secondary. For this demonstration, two CARP VIPs will be added: One for WAN, and one for LAN.

1. Navigate to **Firewall > Virtual IPs** on the primary node.
2. Click  at the top of the list to create a new VIP
3. Set **Type** to *CARP*
4. Set **Interface** to *WAN*
5. Enter the WAN CARP VIP into the **IP Address(es)** section **Address** box and pick the appropriate subnet mask. For this example, enter `198.51.100.200` and `24` (See *WAN IP Address Assignments*).
6. Enter a random password in **Virtual IP Password**. This need only match between the two nodes, which will be handled by synchronization.
7. Select an unused **VHID Group** as determined in *Determine CARP VHID Availability*.

   Note: A common tactic is to make the VHID match the last octet of the IP address, so in this case `200` would be chosen.

8. Set the **Advertising Frequency** to a **Base** of *1* and a **Skew** of *0*. This value will be automatically adjusted when it is copied to the secondary.
9. Enter a **Description** such as `WAN CARP VIP`.
10. Click **Save**
11. Click **Apply Changes**

The **Base** and **Skew** together determine how often a CARP heartbeat is sent. The value of **Base** adds whole seconds and should match between the two nodes. The **Skew** value adds 1/256th of a second increments. The primary node should always have a **Skew** of 0 or 1. The secondary node must be higher, typically 100+. That adjustment is handled automatically by the configuration synchronization process.

Note: If CARP appears to be too sensitive to latency on a given network, adjusting the **Base** by adding one second at a time is recommended until stability is achieved.

Repeat the above process for the `LAN CARP VIP`:

1. Navigate to **Firewall > Virtual IPs**
2. Click Add at the top of the list to create a new VIP
3. Set **Type** to *CARP*
4. Set **Interface** to *LAN*
5. Enter the LAN CARP VIP into the **IP Address(es)** section **Address** box and pick the appropriate subnet mask. For this example, enter `192.168.1.1` and `24` (See *LAN IP Address Assignments*).
6. Enter a random password in **Virtual IP Password**.
7. Select an unused **VHID Group** as determined in *Determine CARP VHID Availability*.
8. Set the **Advertising Frequency** to a **Base** of *1* and a **Skew** of *0*.
9. Enter a **Description** such as `LAN CARP VIP`.
10. Click **Save**
11. Click **Apply Changes**

If there are any additional IP addresses in the WAN subnet that will be used for purposes such as 1:1 NAT, port forwards, VPNs, etc, they may be added now as well.

Check **Firewall > Virtual IPs** on the secondary node to ensure that the VIPs synchronized as expected.

The Virtual IP addresses on both nodes will look like the following if the process was successful.

Virtual IP Address

Virtual IP address	Interface	Type	Description	Actions
198.51.100.200/24 (vhid: 200)	WAN	CARP	WAN CARP VIP	
192.168.1.1/24 (vhid: 1)	LAN	CARP	LAN CARP VIP	

Add

Fig. 2: CARP Virtual IP Address List

Check CARP Status

Now visit **Status > CARP** on both nodes to confirm the proper status. The primary node should indicate MASTER status for all VIPs, and the secondary node should indicate BACKUP status for all VIPs. If the status is incorrect, see *Troubleshooting High Availability*.

CARP Interfaces

CARP Interface	Virtual IP	Status
WAN@200	198.51.100.200/24	MASTER
LAN@1	192.168.1.1/24	MASTER

Fig. 3: CARP VIP Status on Primary

CARP Interfaces

CARP Interface	Virtual IP	Status
WAN@200	198.51.100.200/24	BACKUP
LAN@1	192.168.1.1/24	BACKUP

Fig. 4: CARP VIP Status on Secondary

8.2.5 Setup Manual Outbound NAT

Now it is time to put the new CARP VIPs to use. The NAT settings will synchronize so these changes need only be made to the primary node.

1. Navigate to **Firewall > NAT**, **Outbound** tab on the primary node
2. Change the **Mode** to *Manual Outbound NAT rule generation*
3. Click **Save**, the rule list will be populated with rules equivalent to what was in use for the default, Automatic Outbound NAT.

   Note: If no rules appear in the list, ensure the WAN has a gateway selected under **Interfaces > WAN**

4. Click  to edit the rule for the LAN subnet
5. Set **Translation** to the WAN CARP VIP, *198.51.100.200* in this example.
6. Click **Save**
7. Repeat that edit for each rule in the list except the rules with a source of *127.0.0.0/8*.
8. Click **Apply Changes**
9. Visit **Firewall > NAT**, **Outbound** tab on the secondary node to ensure the rule changes are reflected there.

☐	✔	WAN	192.168.1.0/24	*	*	500	198.51.100.200	*	✔	Auto created rule for ISAKMP - LAN to WAN
☐	✔	WAN	192.168.1.0/24	*	*	*	198.51.100.200	*	⇄	Auto created rule - LAN to WAN

Fig. 5: Outbound NAT Rules for LAN with CARP VIP

Warning: If additional local interfaces are added later, such as a second LAN, DMZ, etc, and that interface uses private IP addresses, then additional manual outbound NAT rules must be added at that time.

8.2.6 Other NAT Concerns

If there are any port forwards to be added using the WAN CARP VIP, they may be added now using **Firewall > NAT**, **Port Forward** tab. Port forwards will work the same as usual, but the **Destination** will be the WAN CARP VIP.

8.2.7 Setup DHCP

The DHCP server daemons on the cluster nodes need adjustments so that they can work together. The changes will synchronize from the primary to the secondary, so as with the VIPs and Outbound NAT, these changes need only be made on the primary node.

1. Navigate to **Services > DHCP Server**, **LAN*** tab.
2. Set the **DNS Server** to the LAN CARP VIP, here `192.168.1.1`
3. Set the **Gateway** to the LAN CARP VIP, here `192.168.1.1`
4. Set the **Failover Peer IP** to the actual LAN IP address of the secondary node, here `192.168.1.3`
5. Click **Save**

Setting the **DNS Server** and **Gateway** to a CARP VIP ensures that the local clients are talking to the failover address and not directly to either node. This way if the primary fails, the local clients will continue talking to the secondary node.

The **Failover Peer IP** allows the daemon to communicate with the peer directly in this subnet to exchange data such as lease information. When the settings synchronize to the secondary, this value is adjusted automatically so the secondary points back to the primary.

On both nodes, visit **Status > DHCP Leases** to see the status. A section will be displayed at the top containing the failover pool status, one line will be shown for each local interface pool. When the two nodes are working properly, both will indicate a "normal" status.

Failover Group	My State	Since	Peer State	Since
dhcp_lan (LAN)	normal	2015/06/25 15:55:55	normal	2015/06/25 15:55:55

Fig. 6: DHCP Failover Status

8.2.8 VPNs and Other Services

When configuring a VPN, such as OpenVPN or IPsec, pick a WAN CARP VIP as the **Interface** for the VPN and ensure the remote peer also builds the other side of the tunnel using the CARP VIP as the peer address.

For other local services, packages, etc. likewise a CARP VIP is recommended for binding if the service will work on both nodes.

High Availability support in packages varies. Check the package documentation for information on if, or how, various aspects of High Availability work with a specific package.

8.2.9 Additional Interfaces

Additional local interfaces may also be configured, repeating some of the previous steps as needed:

1. Assign the interface on both nodes identically
2. Enable the interface on both nodes, using different IP addresses within the same subnet
3. Add a CARP VIP inside the new subnet (Primary node only)
4. Add firewall rules (Primary node only)
5. Add Manual Outbound NAT for a source of the new subnet, utilizing the CARP VIP for translation (Primary node only)

6. Configure the DHCP server for the new subnet, utilizing the CARP VIP for DNS and Gateway roles (Optional, Primary node only)

8.3 Components of a High Availability Cluster

Though often erroneously called a “CARP Cluster”, two or more redundant pfSense firewalls are more aptly titled a “High Availability Cluster”, since CARP is only one of several technologies used to achieve High Availability with pfSense.

The most common High Availability cluster configuration includes only two nodes. It is possible to have more nodes in a cluster, but they do not provide a significant advantage. This guide assumes two SG-4860 nodes are in use, one acting as the primary node and the other as the secondary node.

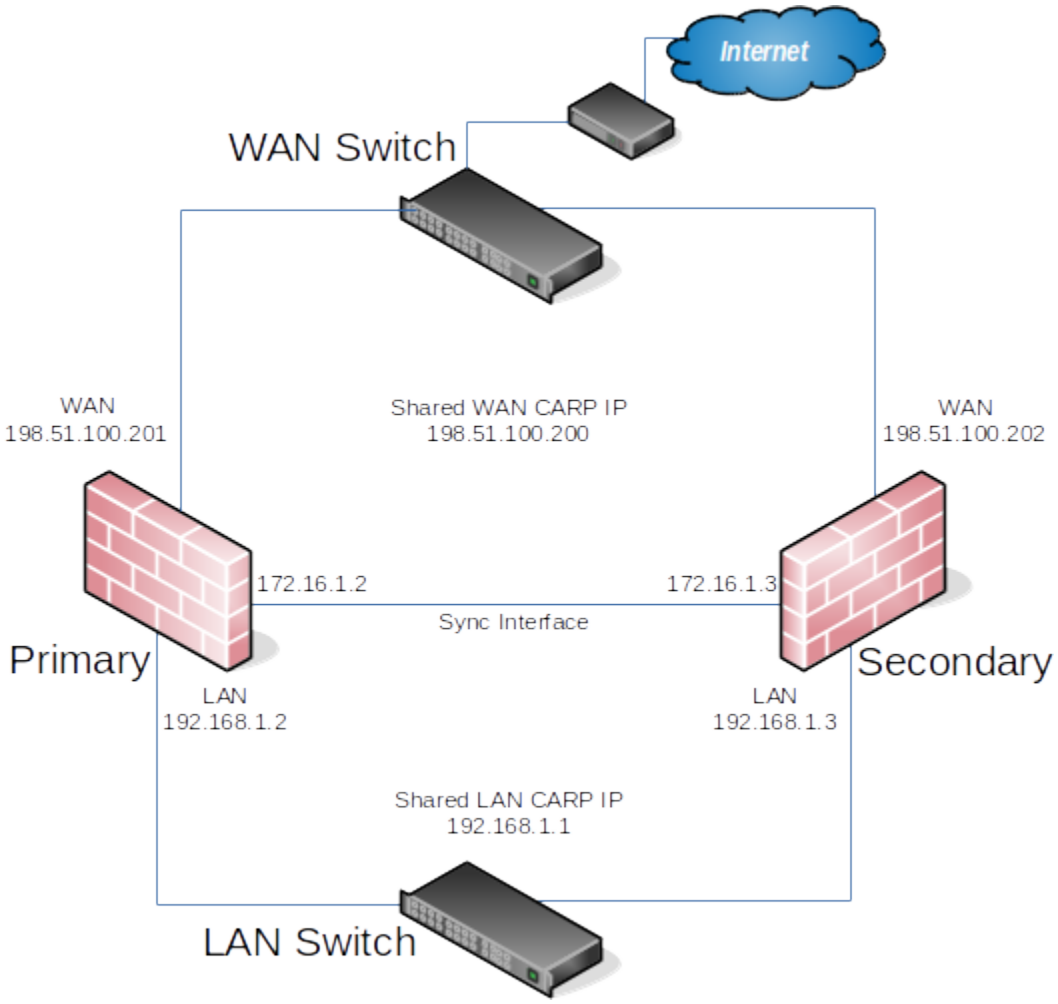

Fig. 7: Example High Availability Cluster

8.3.1 IP Address Redundancy (CARP)

For connectivity through a cluster to continue seamlessly during failover, traffic to and from the cluster must use redundant IP addresses. In pfSense, Common Address Redundancy Protocol (CARP) is used for this task.

A CARP type Virtual IP address (VIP) is shared between nodes of a cluster. One node is master and receives traffic for the IP address, and the other nodes maintain backup status and monitor for heartbeats to see if they need to assume the master role if the previous master fails. Since only one member of the cluster at a time is using the IP address, there is no IP address conflict for CARP VIPs.

In order for failover to work properly it is important that inbound traffic coming to the cluster (routed upstream traffic, VPNs, NAT, local client gateway, DNS requests, etc) be sent to a CARP VIP and for outgoing traffic such as Outbound NAT to be sent from a CARP VIP. If traffic is addressed to a node directly and not a CARP VIP, then that traffic will not be picked up by other nodes.

CARP works similar to VRRP and HSRP, and may even conflict in some cases. Heartbeats are sent out on each interface containing a CARP VIP, one heartbeat per VIP per interface. At the default values for skew and base, a VIP sends out heartbeats about once per second. The skew determines which node is master at a given point in time. Whichever node transmits heartbeats the fastest assumes the master role. A higher skew value causes heartbeats to be transmitted with more delay, so a node with a lower skew will be the master unless a network or other issue causes the heartbeats to be delayed or lost.

Note: Never access the firewall GUI, SSH, or other management mechanism using a CARP VIP. For management purposes, only use the actual IP address on the interface and not the VIP. Otherwise it cannot be determined beforehand which unit is being accessed.

8.3.2 Configuration Synchronization (XMLRPC)

To make the job of maintaining two practically identical firewall nodes easier, configuration synchronization is possible using XMLRPC. Some areas cannot be synchronized, such as the Interface configuration, but many other areas can: Firewall rules, aliases, users, certificates, VPNs, DHCP, routes, gateways, and more. As a general rule, items specific to hardware or a particular installation, such as Interfaces or values under **System > General** or **System > Advanced** do not synchronize. For a list of areas that will synchronize, see the checkbox items on **System > High Avail Sync** in the XMLRPC section. Most packages will not synchronize but some contain their own synchronization settings. Consult package documentation for more details.

When XMLRPC Synchronization is enabled, settings from supported areas are copied to the secondary and activated after each configuration change.

XMLRPC Synchronization is optional, but maintaining a cluster is a lot more work without it.

8.3.3 State Table Synchronization (pfsync)

Active connection state information is exchanged between nodes using the pfsync protocol to achieve seamless failover. When pfsync is active and properly configured, all nodes will have knowledge of each connection flowing through the cluster. If the master node fails, the backup node will take over and clients will not notice the transition since both nodes knew about the connection beforehand.

Failover can still operate without pfsync, but it will not be seamless. Without pfsync if a node fails and another takes over, user connections would be dropped. Users may immediately reconnect through the other node, but they would be disrupted during the transition.

8.4 Testing High Availability

With all of the configuration complete, the time has come for testing. Tests for each aspect of the system are listed below. If any of the tests fails, review the configuration and consult *Troubleshooting High Availability* for assistance.

8.4.1 Verify General Functionality

Setup a client on the LAN and ensure that it receives a DHCP IP address and that it shows the LAN CARP VIP as its gateway and DNS server. Verify that the client can reach the Internet and otherwise function as expected.

8.4.2 Verify XMLRPC Sync is working

XMLRPC Configuration Synchronization can be tested several ways. The easiest method is to make a change to any supported area on the primary, such as a firewall rule, and then see if the change is reflected on the secondary after a few moments.

The manual method for forcing a synchronization task to test XMLRPC is to visit **Status > Filter Reload** on the primary node and click **Force Config Sync**. The status will change briefly and then if everything is working properly, a message will be displayed indicating the sync completed successfully.

8.4.3 Verify CARP is working

Visit **Status > CARP** on both nodes to check if CARP is functional. The primary node will display "MASTER" for all CARP VIPs and the secondary will display "BACKUP" for all CARP VIPs. If the status screen indicates that CARP is disabled, press the **Enable CARP** button.

8.4.4 Verify State Synchronization is working

The **Status > CARP** page lists **pfsync nodes** which give an indication of the state synchronization status. The values may not always match identically on both nodes, but they will be close. If the two are very different, it can indicate a problem with state synchronization. If they are identical or nearly identical, then state synchronization is working.

8.4.5 Testing Failover

A manual failover test may be initiated in one of four ways:

1. Click **Temporarily Disable CARP** on **Status > CARP** on the primary node. This will disable CARP temporarily, and if the primary node is rebooted it will turn back on. Click **Enable CARP** to turn it back on.
2. Click **Enter Persistent CARP Maintenance Mode** on **Status > CARP** on the primary node. This will disable CARP persistently, even if the primary node is rebooted. To exit maintenance mode, click **Leave Persistent CARP Maintenance Mode** to enable CARP once again.
3. Unplug a network cable from an interface with a CARP VIP present, such as WAN or LAN. This will trigger a failover event. Plug the cable back in to recover.
4. Shut down or reboot the primary node.

During any of the above tests, visit **Status > CARP** on the secondary to confirm that the CARP VIPs have taken over and show a "MASTER" status.

Before, during, and after triggering a failover, test connections from a client on the LAN through to the Internet to ensure connectivity works at each step. Downloading a file, streaming audio, or streaming video will most likely continue uninterrupted. VoIP-based phone calls may have a slight disruption as they are not buffered like the others.

Also have a client attempt to obtain an IP address by DHCP while running from the secondary.

If VPNs or other services have been configured, check those during the test as well to ensure the VPN established on the secondary node and continues to pass traffic.

Once the primary node has returned to "MASTER" status, ensure everything continues to work.

8.5 Troubleshooting High Availability

In the event that any of the testing fails (*Testing High Availability*), there are a few common things to check.

8.5.1 Review the Configuration

Before digging too deep into the technical details below, first review the configuration and ensure all steps were followed accurately.

8.5.2 Troubleshooting CARP

Check Interface Status

If an interface shows "INIT" for the CARP state, as shown in *CARP Status on Primary with Disconnected Interface*, most commonly this indicates that the interface upon which this VIP resides is not connected to anything. If there is no link to a switch or another port, the interface is down and the VIP cannot be fully initialized. If the NIC is plugged in and appears to have a link when this occurs, edit, save, and apply changes for the VIP in question to reconfigure it.

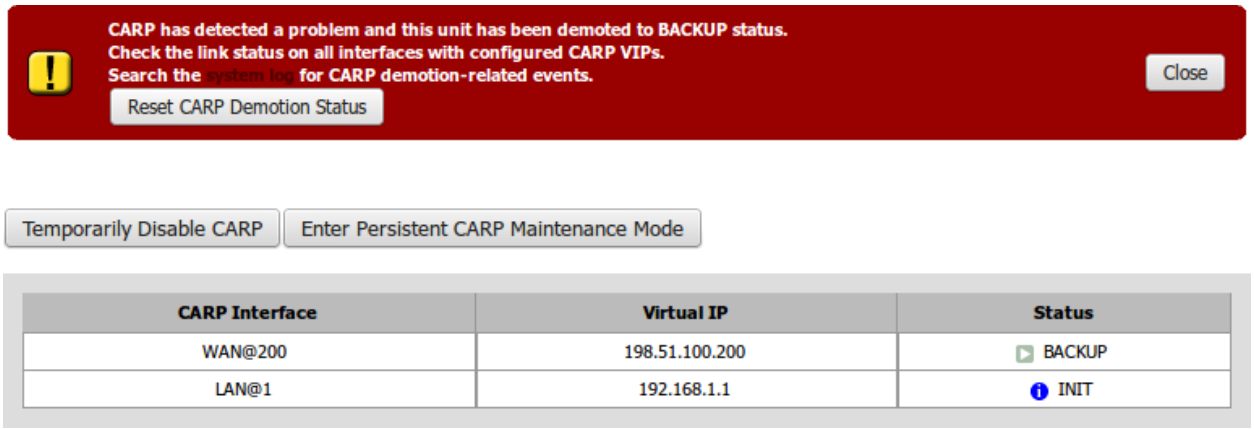

Fig. 8: CARP Status on Primary with Disconnected Interface

Conflicting VHIDs

The VHID determines the virtual MAC address used by that CARP IP. The input validation in pfSense will not permit using conflicting VHIDs on a single pair of systems, however if there are multiple systems on the same broadcast domain running CARP, it is possible to create a conflict. VRRP also uses the same virtual MAC address scheme, so a VRRP IP using the same VRID as a CARP IP VHID will also generate the same MAC address conflict.

When using CARP on the WAN interface, this also means VRRP or CARP used by the ISP can also conflict. Be sure to use VHIDs that are not in use by the ISP on that broadcast domain.

In addition to creating a MAC conflict which can interfere with traffic, it can also interfere with the CARP VIP status.

Incorrect Subnet Mask

The subnet mask for a CARP VIP must match the subnet mask on the Interface IP address for the same subnet. For example, if an interface IP address is 192.168.1.2/24, the CARP VIP must also be 192.168.1.1/24.

Switch/Layer 2 Issues

Typically a switch or layer 2 issue manifests itself as both units showing "MASTER" status for one or more CARP VIPs. If this happens, check the following items:

1. Ensure that the interfaces on both boxes (The WANs, LANs, etc, etc) are connected to the proper switch/VLAN/layer 2. For example, ensure that the LAN port on both units is connected to the same switch/VLAN.
2. Verify that the two nodes can reach each other (via ICMP echo, for example) on each segment. Firewall rules may need to be added to WAN to accommodate this test.
3. If the units are plugged into separate switches, ensure that the switches are properly trunking and passing broadcast/multicast traffic.
4. If the switch on the back of a modem/CPE is being used, try a real switch instead. These built-in switches often do not properly handle CARP traffic. Often plugging the firewalls into a proper switch and then uplinking to the CPE will eliminate problems.
5. Disable IGMP snooping or other multicast limiting and inspecting features. If they are already off, try enabling the feature and disabling it again.

8.5.3 Troubleshooting XMLRPC

If an XMLRPC synchronization attempt fails, a notice is generated in the GUI to bring attention to it, as seen in *XMLRPC Sync Failure Notice*. The notice typically contains some information about why it failed that points to a fix, but if that is not enough, check the other items in this section.

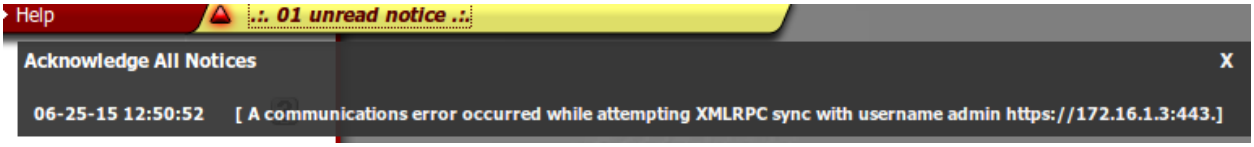

Fig. 9: XMLRPC Sync Failure Notice

Check the System Log

XMLRPC failure details are logged to the main system log (**Status > System Logs**, **General** tab). Usually the error is stated plainly, for example an authentication failure would indicate that the password entered for the Admin user on the synchronization settings was incorrect. As shown in *XMLRPC Sync Failure Log Entry* a timeout happened during the synchronization attempt. In this example it was due to a missing firewall rule.

Jun 25 12:50:52	php-fpm[244]: /rc.filter_synchronize: XML_RPC_Client: Connection to RPC server 172.16.1.3:443 failed. Operation timed out 103
Jun 25 12:50:52	php-fpm[244]: /rc.filter_synchronize: A communications error occurred while attempting XMLRPC sync with username admin https://172.16.1.3:443.

Fig. 10: XMLRPC Sync Failure Log Entry

Check the Firewall Log

Visit **Status > System Logs**, **Firewall** tab on the secondary node. Check the log for entries failing to reach the secondary's Sync interface on the GUI port, as seen in *XMLRPC Sync Failure Firewall Log Entry*. If the traffic is shown as blocked, adjust the Sync interface rules as needed.

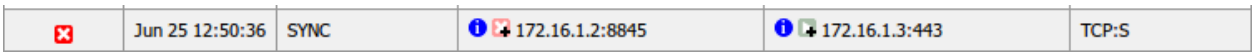

Fig. 11: XMLRPC Sync Failure Firewall Log Entry

Check the Admin User

Visit **System > User Manager** and ensure that the admin user is enabled on both systems and that the admin password is the same on both systems. Visit **System > High Avail Sync** and double check that the admin username has been entered and that the correct password is present.

Verify Connectivity

Check **Status > Interfaces** and ensure the Sync interface shows a link on both units. If there is no link, ensure a cable is connected between the two units. The ports on the SG-4860 are Auto-MDIX so either a straight-through patch or a crossover cable will work. If a short cable is in use, try a longer cable (minimum 3ft/1m). If a link can still not be achieved, try using a small switch or VLAN between the two nodes.

Add a firewall rule to the Sync interface to allow ICMP echo requests and then attempt to ping from one firewall to the other to ensure they can reach each other at layer 3. If they cannot, double check the interface IP address and subnet mask settings, along with the cabling.

8.5.4 Troubleshooting pfsync

If the pfsync nodes do not line up under **Status > CARP**, that can indicate that the states have not been synchronized.

Check Firewall Rules

Check the firewall log at **Status > System Logs**, **Firewall** tab on both nodes. If any pfsync protocol traffic is present, the firewall rules on the Sync interface are probably incorrect.

Look at **Firewall > Rules** on the **Sync** interface tab. Make sure that the rules will pass *pfsync* protocol traffic, or traffic of *any* protocol, to *any* destination. Adjust the rules accordingly and check the logs and CARP status again to see if it starts working.

Verify Connectivity

See *Verify Connectivity* above to check the connection between the nodes.

Check Interfaces

If the states appear to sync but failover is still not seamless, check **Interfaces > (Assign)** and make sure the interfaces all line up physically as well as by name. In pfSense 2.2 and later, the states are bound to the interface so if, for example, the LAN interface is igb0 on one unit but igb3 on the other, then the states will not line up. Fix the interfaces so they are identical on both units.

8.5.5 Troubleshooting Local Services

DNS Resolution

If local clients are unable to obtain DNS responses from a CARP VIP on the cluster, check the following items:

- If using the default DNS Resolver (unbound), visit **Services > DNS Resolver** and click Save on the primary to ensure the default values are fully respected.
- If using either the DNS Resolver or DNS Forwarder, ensure the daemon is configured to listen on All interfaces or at least Localhost and the internal CARP VIPs.
- Ensure the local interface firewall rules pass both TCP and UDP port 53 to the CARP VIPs used for local DNS.
- Ensure the firewall itself has DNS servers configured under **System > General**, especially if using the DNS Forwarder (dnsmasq) instead of the DNS Resolver (unbound).

DHCP

If the DHCP failover pool status does not reach "normal", there are a few items to check:

- Ensure both units are connected to the same switch/subnet on the correct interface.
- Verify connectivity between the two units on that interface.
- Ensure the failover peer IP address has been properly configured
- Ensure that there is a CARP VIP on the interface in question
- Ensure that the CARP VIP on the primary node has a skew of 0 or 1, and the secondary has a skew of 100 or higher.
- If all else fails:
  - Click  to stop the DHCP service from **Status > Services** on both nodes
  - Visit **Diagnostics > Command** Prompt on both nodes
  - Run the following command in the **Shell Execute** box on both nodes: `rm /var/dhcpd/var/db/dhcpd.leases*`
  - Click  to start the DHCP service from **Status > Services** on both nodes

8.6 Upgrading pfSense on a High Availability Cluster

There is more to updating a cluster than the typical process, but in all updating a cluster is much less disruptive as the users will not have any downtime in most all cases.

If at any point in this procedure a failure condition is encountered, seek assistance from support.

8.6.1 Review the Changelog and Upgrade Guide

Before starting any part of an upgrade, first look at the Netgate Blog and release changelogs for any notable changes or items to be aware of between the version currently in use and the one that will be in use after upgrading.

Common issues are also listed in the upgrade guide, especially for major version upgrades.

8.6.2 Backup

Before starting, take a fresh backup from **Diagnostics > Backup/Restore** on both nodes.

> **Warning:** Do not skip this step! A backup is quick and easy to do, and invaluable to have if the upgrade does not go as expected!

Download installation media for the release currently in use if a reinstall is necessary.

8.6.3 Upgrade Secondary

Perform the OS upgrade on the secondary node first. This way, if the upgrade fails, there is no interruption and if a reinstall is needed, it can be done without worry.

8.6.4 Test Secondary

Once the secondary has booted back up, login and confirm that it is running as expected. If all services are active, the CARP status is OK, and so on then it is time to test. Force a failover from the primary node by placing it into maintenance mode (See *Testing Failover*) and observe what happens on the secondary. If the secondary takes over OK and traffic continues to flow, then it is OK to proceed.

8.6.5 Upgrade Primary

With the primary node in maintenance mode, it is safe to upgrade without additional interference. Initiate the OS upgrade and let the system reboot. Once it has rebooted, confirm that local services are running as expected and then take the node out of maintenance mode.

8.6.6 Test Again

With both units on the current OS and active, run some final tests to ensure that services are operational, traffic is flowing, and that the CARP, DHCP, and other status areas are all running properly.

CHAPTER

NINE

REINSTALLING PFSENSE

1. Registered users can log in to their Portal account page and download the appropriate factory installer image:

   `pfSense-netgate-memstick-XG-7100-2.4.3-RELEASE-p1-amd64.img.gz`

   If you no longer have an active portal subscription, please contact support to re-enable access to the image free of charge.

   Note: The pfSense factory version is the version that is preinstalled on units purchased from Netgate. The factory image is optimally tuned for our hardware and contains some features that cannot be found elsewhere, such as the AWS VPN Wizard.

   ---

2. Write the image to a USB memstick. Locating the image and writing it to a USB memstick is covered in detail under Writing Flash Drives.
3. Connect to the console port of the pfSense device.

   See also:

   Connecting to Console Port Connecting to the console port. Cable is required.
4. Insert the memstick into an open USB port and boot the system.
5. After a minute the pfSense loader menu will be displayed that contains options to **Boot Multi User**, **Boot Single User**, **Escape to loader prompt**, **Reboot**, select a non-default kernel or configure boot options. Either allow the menu to timeout and boot on its own, or press `1` to boot normally. The factory images of pfSense already have appropriate console port options.
6. The installer will automatically launch once the boot process completes and offer the choice of a Quick/Easy Install, Custom Install, and several other options. Select **Quick/Easy Install** and press `Enter`. Another screen will prompt for confirmation. Select **OK** and press `Enter` to continue.
7. pfSense will be installed to the first available disk in the system. If the system contains an optional SSD storage disk, it will be chosen. Otherwise, the onboard eMMC will be used. It will take a couple of minutes to copy all of the files to the target disk. When the files have finished being copied, the installer will prompt to select either the **Embedded Kernel** or **Standard Kernel**. Select the **Embedded Kernel** and press `Enter`.
8. If you are on the factory version, the installer will then prompt to choose the type of system being installed, which pre-configures device-specific defaults. Choose the option that exactly matches the unit being reinstalled. If the model is unknown, check the sticker on the bottom of the unit.
9. The installer will then prompt to Reboot the system. Select **Reboot** and press `Enter`. The system will reboot.
10. Remove the USB drive from the USB port. pfSense will restart automatically. If the USB drive remains attached, the system will boot into the installer again because by default the system firmware is configured so that a device plugged into the USB port will be booted with a higher priority.

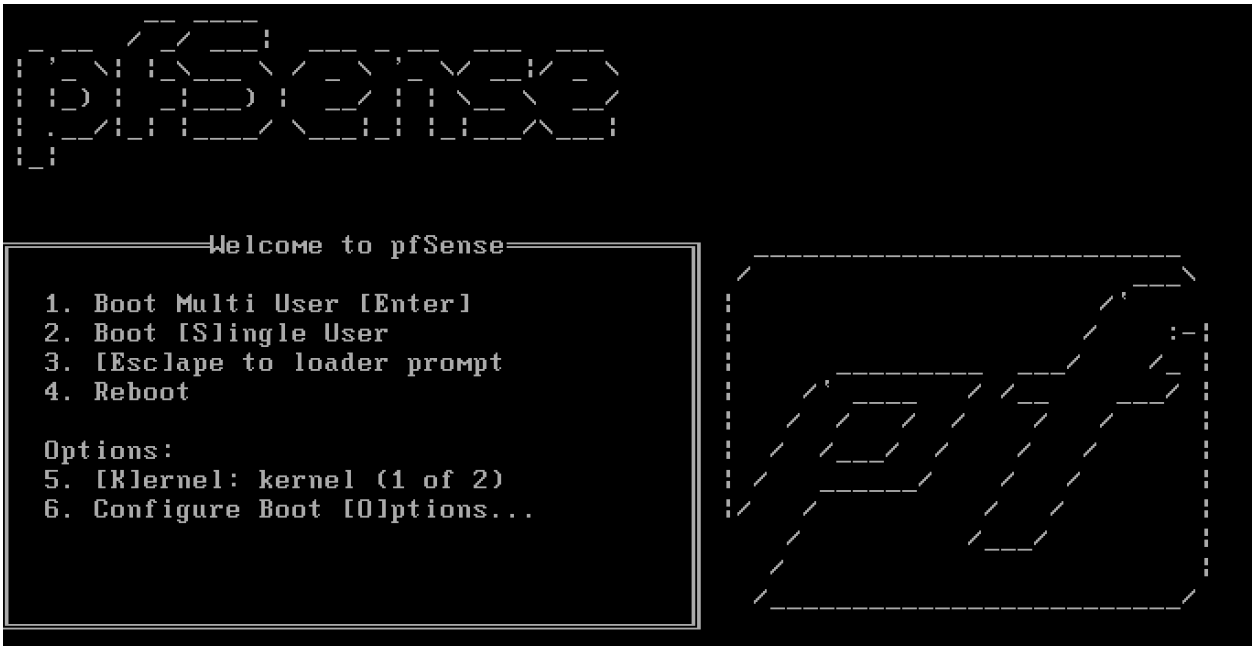


```
pfSense is now rebooting

After the reboot is complete, open a web browser and
enter https://192.168.1.1 (or the LAN IP Address) in the
location bar.

You might need to acknowledge the HTTPS certificate if
your browser reports it as untrusted.  This is normal
as a self-signed certificate is used by default.

*DEFAULT Username*: admin
*DEFAULT Password*: pfsense

Rebooting in 5 seconds. CTRL-C to abort.
Rebooting in 4 seconds. CTRL-C to abort.
Rebooting in 3 seconds. CTRL-C to abort.
Rebooting in 2 seconds. CTRL-C to abort.
Rebooting in 1 second.. CTRL-C to abort.

pfSense is now rebooting.

Dec 14 18:25:18 lighttpd[38936]: (server.c.1567) server stopped by UID = 0 PID =
 35875
```

CHAPTER

TEN

BIOS FLASH PROCEDURE

10.1 Update via the GUI

Warning: This only works with Netgate systems running pfSense version 2.3 or greater.

1. To install the package, navigate to **System -> Package Manager -> Available Packages**.
2. Click the **Install** button for the package named **Netgate_Coreboot_Upgrade**.
3. On the next page, click the **Confirm** button.
4. When the installation is complete a message will appear saying:

```
pfSense-pkg-Netgate_Coreboot_Upgrade installation successfully completed
```

5. Now that the package is installed, navigate to **System -> Netgate Coreboot Upgrade**.
6. This page will show you the latest version of Coreboot available and the current version that is running on the system. If you happen to be on an older version of Coreboot then an **Update** button will be available to click.

Warning: Pay close attention to any disclaimers presented. Some devices require a physical reboot or some step unique to that device.