



Palo Alto Networks Firewall Essentials Installation and Configuration Guide

The Palo Alto Networks Academy Firewall Essentials lab set is designed to have Internet access. Due to this requirement, 2 topologies are needed. The Firewall Essentials Gateway pod (GW) is designed to provide Internet access to underlying Firewall Essentials pods (FE) per host.

This guide includes installation instructions for both the GW pod and the FE pod.

Document Version: 2016-07-21

Copyright © 2016 Network Development Group, Inc.
www.netdevgroup.com

NETLAB Academy Edition, NETLAB Professional Edition, and NETLAB+ are registered trademarks of Network Development Group, Inc.

Palo Alto Networks and the Palo Alto Networks logo are trademarks or registered trademarks of Palo Alto Networks, Inc.

Contents

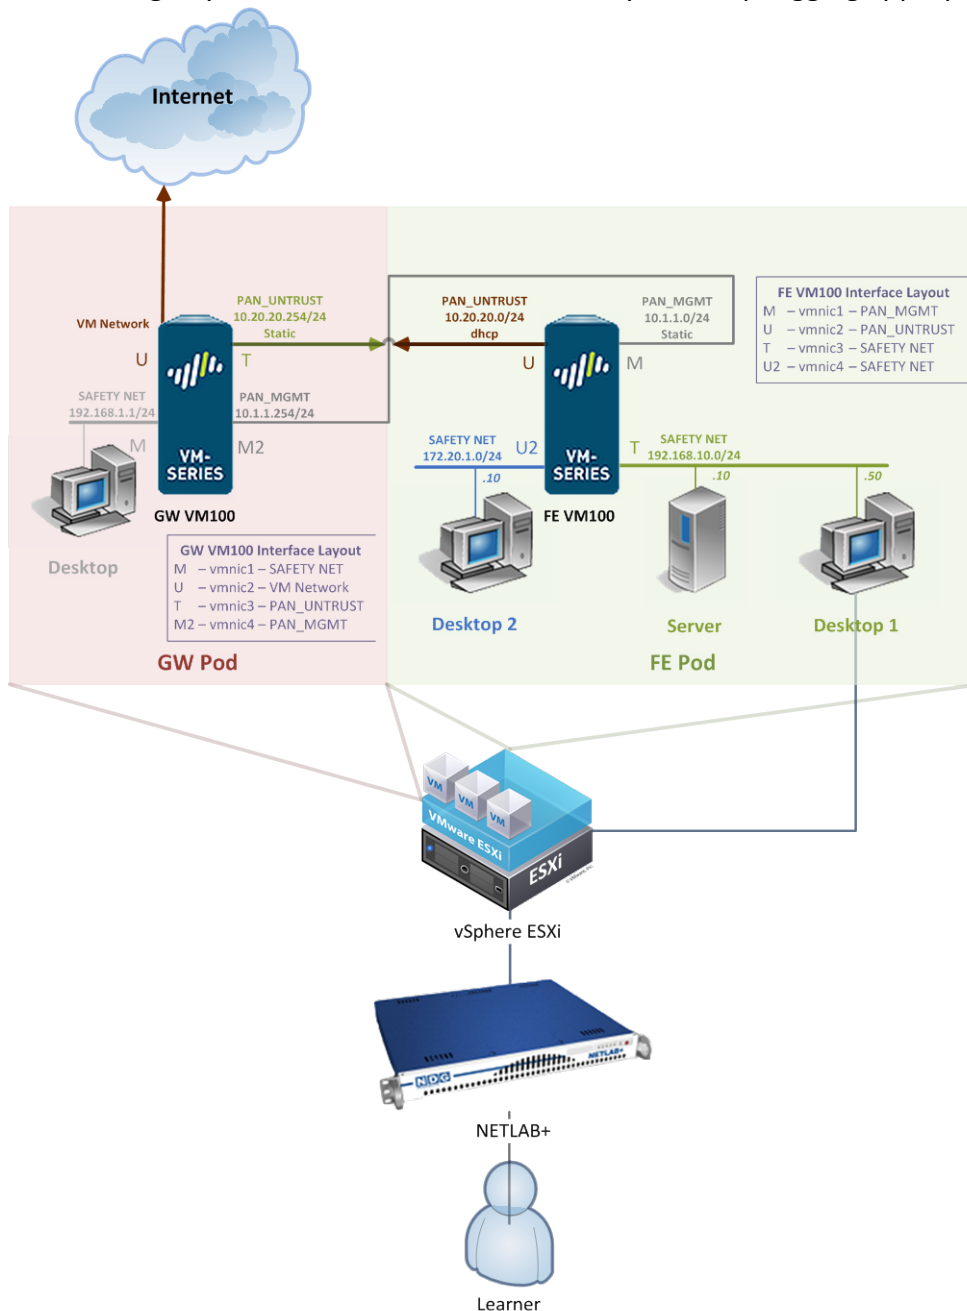
1	Introduction	3
1.1	NETLAB+ Pod Internet Access and Use Agreement	4
1.2	Pod Setup Overview	4
2	Planning.....	6
2.1	Environment.....	6
2.2	Pod Creation Workflow.....	7
2.3	Pod Resource Requirements.....	8
2.4	ESXi Host Server Requirements.....	8
2.5	NETLAB+ Requirements	8
2.6	Software Requirements	9
2.7	Networking Requirements	9
3	Obtaining Software and Licenses.....	10
3.1	Downloading OVF Files.....	10
3.2	Obtaining Software Licenses	10
4	Master Pod Configuration.....	11
4.1	Host Configuration	11
4.1.1	Port Group Configuration	11
4.1.2	NETLAB+ Virtual Machine Infrastructure Setup	14
4.2	Gateway Master (GW) Pod Setup	14
4.2.1	Deploying GW Virtual Machine OVF/OVA Files.....	15
4.2.2	Create Snapshots on the Master Virtual Machines.....	16
4.2.3	NETLAB+ Virtual Machine Inventory Setup	17
4.2.4	Install the Master GW Pod.....	18
4.2.5	Update the Master Pod	19
4.2.6	Bring the GW Master Pod Online	20
4.3	Firewall Essentials Master (FE) Pod Setup	21
4.3.1	Deploying FE Virtual Machine OVF/OVA Files	21
4.3.4	Install the Master FE pod	24
5	Pod Cloning and Configuration	27
5.1	Pod Cloning	27
5.1.1	Linked Clones and Full Clones.....	27
5.1.2	Creating User Pods.....	27
5.2	GW Pod Configuration	29
5.2.1	IP Address Assignment.....	29
5.2.1.1	Static IP Address	30
5.2.1.2	DHCP IP Address	31
5.2.2	DNS Settings.....	31
5.2.3	Licensing.....	33
5.2.4	Startup and Shutdown the Firewall	33
5.3	FE Pod Configuration.....	34
5.3.1	IP Addressing.....	35
5.3.1.1	Boot FE Firewalls - Manual Method	35
5.3.1.2	Boot FE Firewalls - PowerCLI Method	36
5.3.2	Licensing.....	37
5.3.2.1	Troubleshooting.....	39

- 5.3.3 Pod Snapshots..... 40
 - 5.3.3.1 Snapshot the Virtual Machines - Manual Method 41
 - 5.3.3.2 Snapshot the Virtual Machines - PowerCLI Method 42
- 5.4 Bring Pods Online 43
- 6 PAN Firewall Administration Best Practices 44
 - 6.1 Administration..... 44
 - 6.2 Security Policies..... 44
 - 6.3 Logging 44
 - 6.4 Threat Prevention 45
 - 6.4.1 URL Filtering..... 45
 - 6.4.2 Wildfire 45
 - 6.4.3 Monitoring 45

1 Introduction

The Palo Alto Networks Firewall Essentials lab set is required, and thus designed, to have Internet access. Due to this requirement, the use of the lab set requires two pods, one to provide Internet access to pods on the host and the other to clone learner pods from.

You specifically agree to log all Internet usage by users (trainees) made through the Palo Alto Network Academy lab environment, following logging instructions and advice provided by Palo Alto Networks, subject to your compliance with all applicable laws. Note that, because of the nature of lab setup as shown below, you will not be able to track Internet usage by MAC address, so it is vital that you set up logging appropriately.



You agree that you are fully responsible for, and that NDG will have no liability or responsibility for: (a) any Internet use by any users of the Palo Alto Networks Academy lab training environment or any additional lab environments that you set up using Palo Alto Networks firewalls, and (b) monitoring, securing and logging Internet activity occurring through the Palo Alto Networks Academy lab training environment.

IMPORTANT: If you decide to add optional functionality to allow trainees (including without limitation remote trainees) to access and use the Internet through the Palo Alto Networks Academy lab environment, you are solely responsible for configuring and managing the Palo Alto Networks firewalls and associated software that is provided by Palo Alto Networks for Internet access, including without limitation all security features and policies associated with the Palo Alto Networks firewalls.

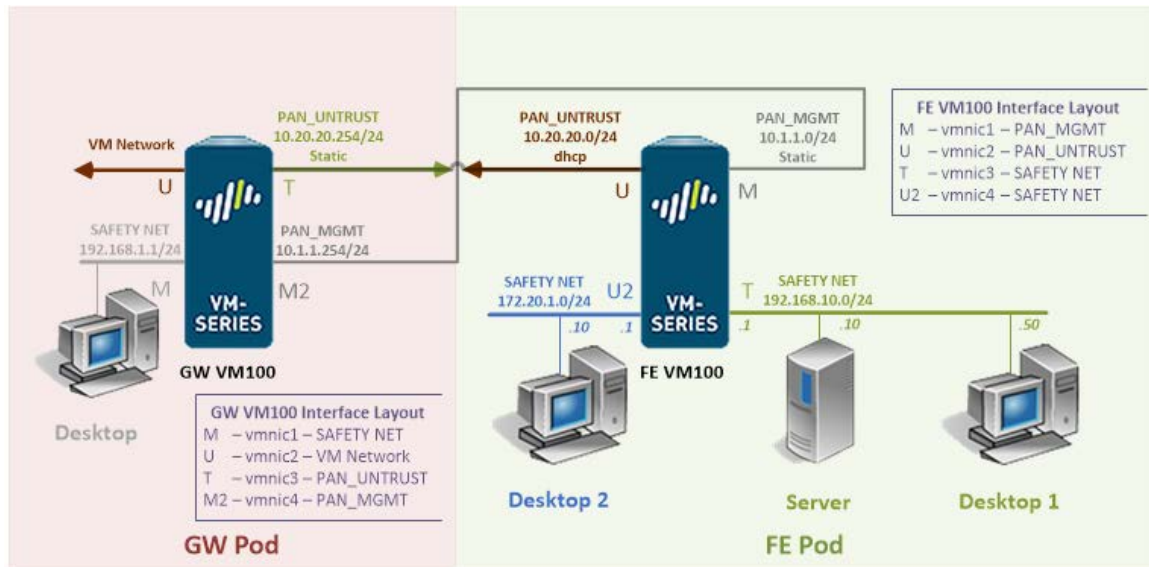
1.1 NETLAB+ Pod Internet Access and Use Agreement

You are required to indicate your acceptance of the NETLAB+ Pod Internet Access and Use Agreement by completing the form at the link below. Your system will not be enabled to support Palo Alto Networks Firewall Essentials pods until the agreement is accepted:

<https://www.netdevgroup.com/content/paloalto/agreement>

1.2 Pod Setup Overview

The Gateway pod (GW Pod) is designed to provide Internet access to underlying Firewall Essentials pods (FE Pod) per host.



Each ESXi host will need special port groups created named PAN_MGMT and PAN_UNTRUST. Then, a single instance of the GW Pod will be deployed on each host that will run the PAN7 FE pods.

The network labeled “VM Network” in the diagram needs to be setup or linked to a port group that has Internet access. A working and routable IP address, static or DHCP assigned, will need to be allocated to vmnic2 of the GW Firewall for the Firewall to communicate out to the Internet.

The PAN_MGMT and PAN_UNTRUST networks are required for the FE Firewall to communicate to the GW Firewall properly. The PAN_UNTRUST on the FE Firewall, identified as interface U in the diagram, is setup to obtain an IP address via DHCP from the GW Firewall T interface.

2 Planning

This guide provides specific information pertinent to delivering the Palo Alto Networks Firewall Essentials course via NETLAB+. It is assumed that you have knowledge of the following prior to attempting deployment of this lab set on your VMware and NETLAB+ infrastructure:

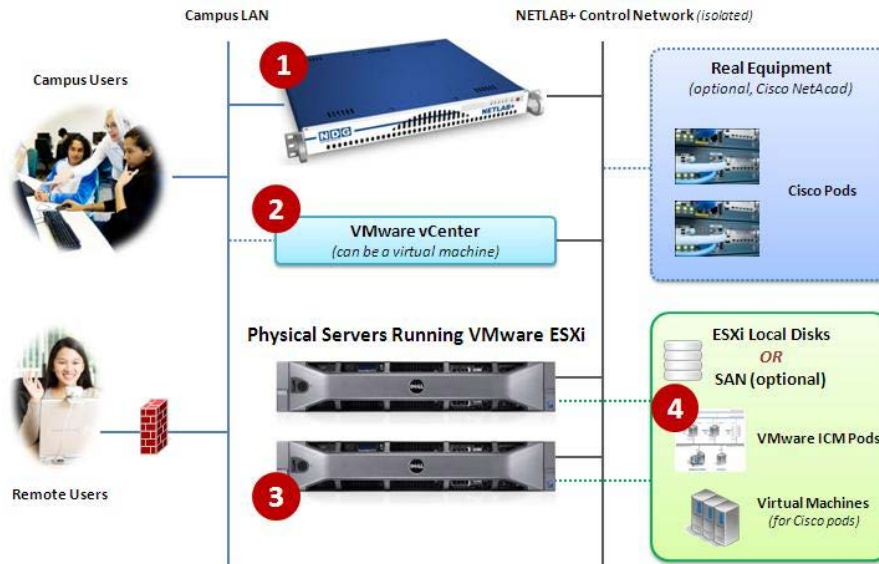
- An understanding and working knowledge of VMware vSphere products and NETLAB+.
- Deploying virtual machines on ESXi.
- Configuring virtual networking in the ESXi environment.
- Virtual machine and virtual pod management concepts using NETLAB+.

Documentation of these topics and more can be found at our website:

<https://www.netdevgroup.com/support/documentation/>

2.1 Environment

The following diagram depicts four major components that make up the training environment.



1. The NETLAB+ server provides the user interface for student and instructor access, an interface to manage virtual machines, and software features to automate pod creation. This document assumes you have already setup your NETLAB+ server.
2. VMware vCenter is used to manage your physical VMware ESXi servers, to create virtual machines, and to take snapshots of virtual machines. NETLAB+ communicates with vCenter to perform automated tasks and

- virtual machine management.
- 3. Physical VMware ESXi servers host the virtual machines in your pods.
- 4. The Palo Alto Networks Firewall Essentials pod consists of 4 virtual machines that reside on your ESXi host(s).

2.2 Pod Creation Workflow

The following list is an overview of the pod setup process.

1. Obtain the master virtual machine images required for the pod.
2. Deploy the master virtual machine images to a master pod.
 - a. Deploy virtual machines using Thin Provisioning to reduce storage consumption.
 - b. Make necessary adjustments to each virtual machine in the environment.
3. Import the deployed virtual machines to the NETLAB+ Virtual Machine Inventory.
4. Take a snapshot of each virtual machine in the master pods labeled GOLDEN_MASTER.
5. Assign and configure pod settings for each virtual machine in each pod.
6. Use the NETLAB+ Pod Cloning feature to create student FE pods from the master FE pod.
7. Configure and license the GW Firewall.
8. License the FE Firewall in all FE student pods.
9. Shutdown FE Firewall and take a GOLDEN_MASTER snapshot of all FE student pod virtual machines.

2.3 Pod Resource Requirements

The Palo Alto Networks Firewall Essentials course will consume 65 GB of storage per each user pod instance.

The following table provides details of the storage requirements for each of the virtual machines in the pod(s).

Pod	Virtual Machine	OVF/OVA	Initial Master Pod (Thin Provisioning)
Gateway	GW Firewall	3.7	20
	Desktop	2.5	7.7
Firewall Essentials	FE Firewall	6.2	13.5
	Desktop1	2.6	6.2
	Desktop2	2.1	5.6
	Server	4.3	12
	Total	21.4	65

2.4 ESXi Host Server Requirements

Please refer to the NDG website for specific ESXi host requirements to support virtual machine delivery: <http://www.netdevgroup.com/content/vmita/requirements/>

The number of **active** pods that can be used simultaneously depends on the NETLAB+ product edition, appliance version and number of VMware ESXi host servers meeting the hardware requirements specifications.

For current ESXi server requirements refer to the following URL:

http://www.netdevgroup.com/support/remote_pc.html#vm_host_server_specifications.

2.5 NETLAB+ Requirements

Installation of the pods as described in this guide requires that your NETLAB+ system is equipped with NETLAB+ version 2011.R5 or later.

Previous versions of NETLAB+ do not support the use of VMware ESXi 5.1, or later, on the physical host servers, which is required to support the installation and use of Palo Alto Networks Firewall Essentials pods on NETLAB+.

2.6 Software Requirements

For the purpose of software licensing, each virtual machine is treated as an individual machine, PC or server. Please refer to the specific vendor license agreements (and educational discount programs, if applicable) to determine licensing requirements for your virtual machines' software, operating system and applications.

The following table lists the software or licenses required for virtual machines inside the Palo Alto Networks Firewall Essentials and Gateway pods.

Pod Software Requirements		
Software	Version	Source
GW Firewall	7.0.6	PAN AuthID
FE Firewall	7.0.6	PAN AuthID

2.7 Networking Requirements

To accommodate the movement of large VMs, OVF/OVAs and ISO disk images from one host to another, a Gigabit Ethernet switch is minimum to interconnect your NETLAB+, vCenter Server system and ESXi host systems.

Protocols required for Internet access are as follows:

- ICMP
- DNS
- HTTP
- HTTPS

3 Obtaining Software and Licenses

3.1 Downloading OVF Files

NDG has built the virtual machines, made available as Open Virtualization Format (OVF) or Open Virtualization Archive (OVA) files. These files are available for download from CSSIA.

To request access to the preconfigured virtual machine templates from CSSIA:

1. Go to the CSSIA Resources page: <http://www.cssia.org/cssia-resources.cfm>.
2. Select **VM Image Sharing Agreement – Image Sharing Agreement**.
3. Select **VM Image Sharing Agreement** to open the request form.
4. Complete and submit your access request by following the instructions on the request form.
5. CSSIA will email a link, along with a username and password to access the download server. Access to the download server is provided only to customers who are current with their NETLAB+ support contract and are participants in the appropriate partner programs (i.e. Cisco Networking Academy, VMware IT Academy, and/or EMC Academic Alliance).
6. Once access to the download server has been established, the virtual machines can be deployed directly to the vCenter Server by clicking on File > Deploy OVF Template in the vClient window and copying the link into the location field.
7. The deployment will start after the username and password are entered.
8. Each virtual machine is deployed individually.

3.2 Obtaining Software Licenses

To obtain licensing and access to the Palo Alto Network 7 Firewall Essentials labs, your institution must be a Palo Alto Networks Authorized Academy Center (AAC).

You can find information about the Palo Alto Networks AAC at the following link:

<https://www.paloaltonetworks.com/services/education/authorized-academy-centers.html>.

Once membership in the Palo Alto Networks AAC is approved, you can request licenses for use with your pods.

You will need one license per pod you intend to install. As an example, for 30 student pods split across 2 host servers:

2 GW pods (one for each host server) + 2 FE Master pods (one for each host server) + 30 FE Student pods (15 on each host server) = 34 licenses.

4 Master Pod Configuration

4.1 Host Configuration

Before deploying virtual machines to your host servers, you must first create two new port groups for use exclusively by the PAN7 FE pods. Please take care to name the port groups exactly as they appear in the following section (case sensitive). This will ensure the VMs can connect to the appropriate networks when pods are started up.

Recent tests have shown the Client Integration Plugin packaged with VCSA 6.0 build 3634788 as functional. Previous versions packaged with VCSA 6 or VCSA 5.5 have had mixed results.

Due to OVF/OVA deployment issues with the vCenter Web Client, we chose to use the vSphere Client for vSphere 6 for creating vswitches, port groups and deploying virtual machines.

4.1.1 Port Group Configuration

The port groups to be created follow below:

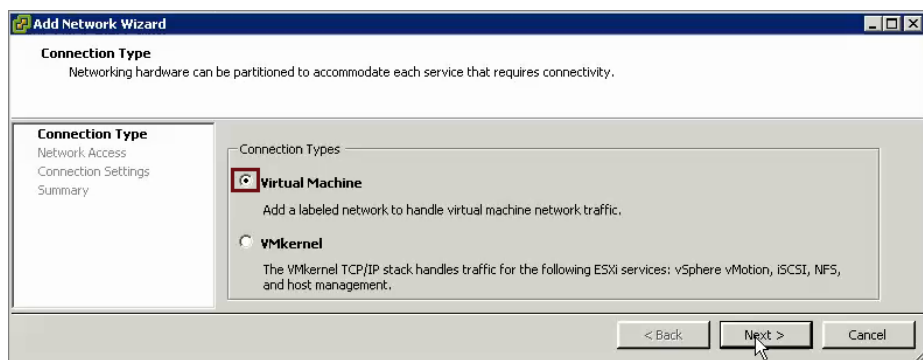
- PAN_MGMT
- PAN_UNTRUST

Use the following instructions to create these port groups.

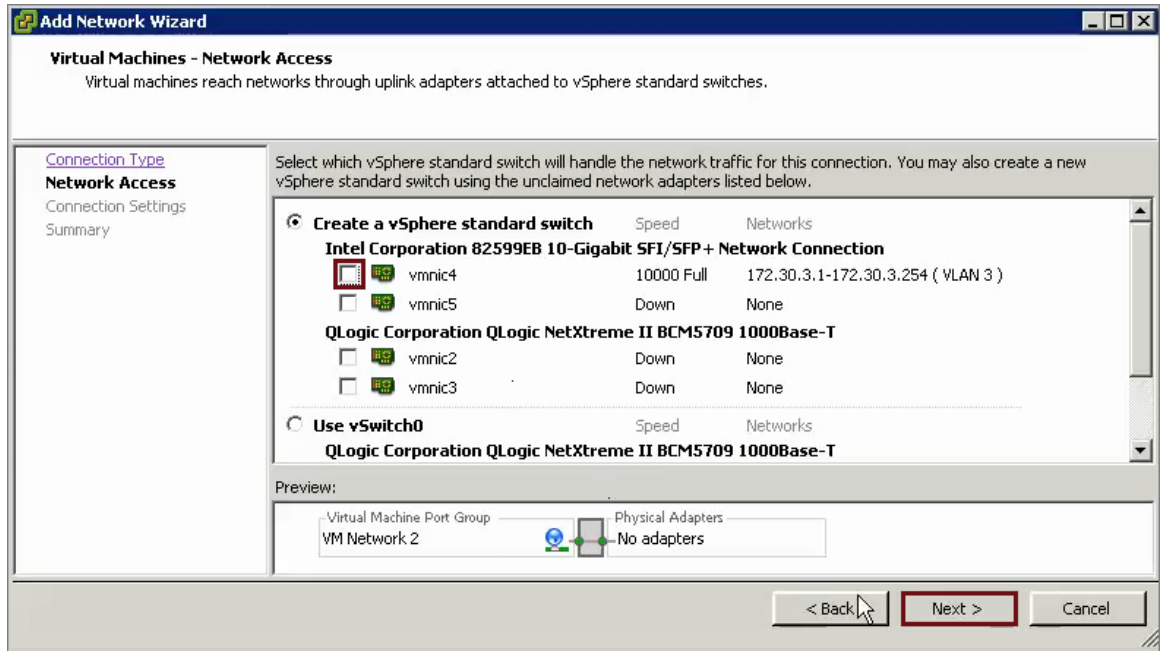
1. Click on the **Add Networking** link under the Configuration tab and Networking section.




2. Select **Virtual Machine** connection type and click **Next**.




- For Network Access, remove the checks from the **checkboxes** by the vmnics to keep the traffic from leaving the virtual switch and then click **Next**.



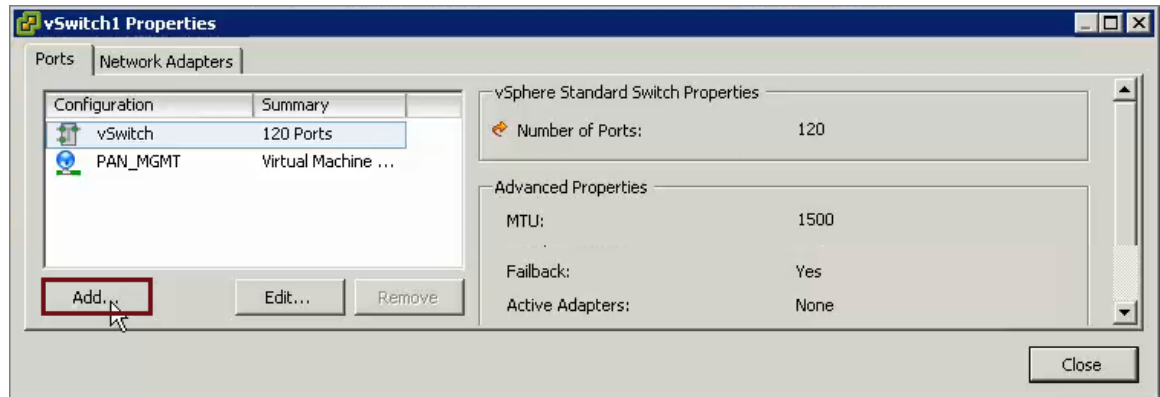
- For Connection Settings, enter **PAN_MGMT** in the **Network Label** and 11 in the **VLAN ID** text boxes and then click **Next**.




- In the Summary, click the **Finish** button.
- Next, click on the **Properties** hyperlink for the vSwitch that you created.



- In the vSwitch Properties window, click on the **Add** button.



- In the Connection Type window, click **Virtual Machine** and then click **Next**.
- In the Connection Settings window, enter **PAN_UNTRUST** in the **Network Label** and **12** in the **VLAN ID** text boxes and then click **Next**.



- Click the Close button to conclude setting up the port groups for the host.

You will need to add these port groups to each host that you intend to run the PAN7 FE pod(s) on. If these networks are not available, the FE pods will not be able to communicate with the GW pod, breaking functionality.

4.1.2 NETLAB+ Virtual Machine Infrastructure Setup

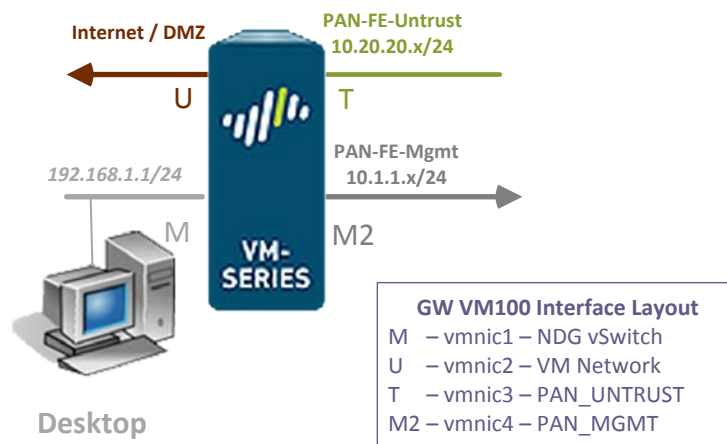
The NETLAB+ Virtual Machine Infrastructure setup is described in the following sections of the: [NETLAB+ Remote PC Guide Series](#).

- *Registering a Virtual Datacenter in NETLAB+*
- *Adding ESXi hosts in NETLAB+*
- *Proactive Resource Awareness*

It is important to configure Proactive Resource Awareness to maximize the number of active pods per physical ESXi host.

4.2 Gateway Master (GW) Pod Setup

The instructions in this sub-section assist you in installing the PAN7 Gateway (GW) pod on your host server. The GW pod is required to be installed on each host server where you plan to deploy the PAN7 FE pod VMs and acts as the gateway for those pods to the internet. You will need to assign an IPv4 address to the Internet / DMZ interface and ensure that the interface is assigned to a port group on a vswitch that has access to that network.



In preparation for the installation of the GW pod, you need the following information:

DHCP or Static IP address configuration?


If you intend to use a static IP, please obtain and record the following information:

IP Address _____
 Subnet Mask _____
 Gateway IP _____

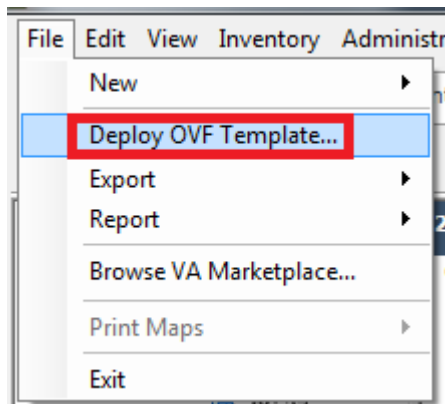
4.2.1 Deploying GW Virtual Machine OVF/OVA Files

The OVF/OVA files that you downloaded from CSSIA must be deployed to your host server.

1. Open the **vClient** on your administration machine, ensure that your downloaded OVA/OVF files are accessible on this machine and then connect to your **vCenter Server**.
2. Select **Hosts and Clusters** in the address bar.



3. Click on the target ESXi Host Server.
4. Click on **File -> Deploy OVF Template**.



5. Click on **Browse** and locate the **PAN7_GW_FM_<VM Name>** OVF/OVA files you downloaded from CSSIA. Click **Next** to continue.
6. On the OVF Template Details window, click **Next**.
7. On the Name and Location window, ensure the deployed name contains **PAN7_GW_Master_<VM Name>** and click **Next**.
8. On the Datastore window, select the target datastore and click **Next**.
9. On the Disk Format window, select **Thin provisioned format** and click **Next**.
10. Perform the following instructions on the Firewall virtual machine only. Leave the default networks on non-Firewall virtual machines.
 - a. On the Gateway Firewall **Network Mapping** window, ensure the following networks are set and then click **Next**.

Virtual NIC	Setting
vmnic1	SAFETY NET
vmnic2	VM Network *
vmnic3	PAN_UNTRUST
vmnic4	PAN_MGMT

* The **VM Network** on vmnic2 will need to be the name of the port group that you intend to use for Internet access. In most cases, **VM Network** is the network used for internet access; however, this is not always the case.

11. On the **Ready to Complete** window, confirm the information and then click **Finish**.
12. vCenter will begin deploying the virtual machine. This may take some time depending on the speed of your connection, HDDs, etc. When completed, click on **Close**.
13. Perform the previous steps for each virtual machine in the GW pod.

4.2.2 Create Snapshots on the Master Virtual Machines

In order to proceed with pod cloning, snapshots must be created for the Master virtual machines.

Verify that all VMs are powered off before taking snapshots.

1. Open the **vClient** on your management workstation. Connect to your **vCenter Server**.
2. Select **Hosts and Clusters** in the address bar.



3. Right-click on each virtual machine and select **Snapshot > Take Snapshot**.
4. Enter **GOLDEN_MASTER** as the Snapshot Name.
5. Enter a description. It is a good idea to include the date in the description for later reference.
6. Click **OK**.
7. Repeat Steps 3-6 for the remaining virtual machines in the pod.
8. When all tasks have completed, log out of the vClient software.

4.2.3 NETLAB+ Virtual Machine Inventory Setup

This section will guide you in adding your templates to the Virtual Machine Inventory of your NETLAB+ system.

1. Login into your NETLAB+ system using the administrator account.
2. Select the Virtual Machine Infrastructure link.



[Virtual Machine Infrastructure](#)


3. Click the Virtual Machine Inventory link.




Virtual Machine Inventory

Import, clone, and manage the inventory of virtual machines to be used with NETLAB+.


4. Click the Import Virtual Machines button.



5. Select the checkbox next to your Palo Alto Networks Firewall Essentials Gateway virtual machines and click Import Selected Virtual Machines.



6. When the Configure Virtual Machines window loads, you can set your virtual machine parameters.
 - a. Check the drop-down box for the correct operating system for each imported virtual machine.
 - b. Select **Master** for the role of each virtual machine.
 - c. Add any comments for each virtual machine in the box to the right.
 - d. Verify your settings and click Import Selected Virtual Machines.



- e. Click OK when the virtual machines have finished loading.
- f. Verify that your virtual machines show up in the inventory.


4.2.4 Install the Master GW Pod

This section will assist you in adding the Palo Alto Networks Firewall Essentials Gateway (GW) pod to your NETLAB+ system.


1. Login into NETLAB+ with the administrator account.
2. Select the **Equipment Pods** link.



3. Create a new pod by scrolling to the bottom and clicking the **Add a Pod** button.



4. On the New Pod Wizard page, click **Next**.
5. Then select the PAN7 Gateway pod radio button and click **Next**.



6. Select a Pod ID and click **Next**.

It is best practice to use a block of sequential ID numbers to number Palo Alto Networks Firewall Essentials Gateway pods you are going to install. The Pod ID number determines the order in which the Palo Alto Networks Firewall Essentials Gateway pods will appear in the scheduler.

7. Type in **PAN7_GW_Master** for the Pod Name and click **Next**.
8. To finalize the wizard click **OK**.

4.2.5 Update the Master Pod


1. Update the Master Pod on your NETLAB+ system.
 - a. Login into NETLAB+ with the administrator account.
 - b. Select the **Equipment Pods** link.




2. Click on the Magnifying Glass icon next to the virtual machine you are about to assign. Please note that your PC IDs will not match the graphic below.

POD 1000 - PCs AND SERVERS (click the GO buttons to reconfigure)						
GO	NAME	PC ID	STATUS	TYPE / VM	OPERATING SYSTEM	
	Desktop	4100	ONLINE	ABSENT		
	Firewall	4101	ONLINE	ABSENT		

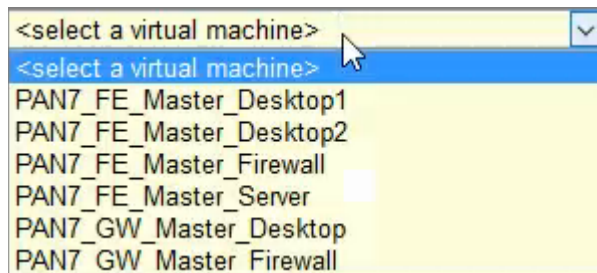
3. Click on **Modify PC Settings**.



4. Change the PC Type drop-down box to **Use Virtual Machine Inventory**.



5. In the Base Virtual Machine window, select your Palo Alto Networks Firewall Essentials Gateway virtual machine to associate with the spot held for it in the pod.



- Update the **Base Snapshot** and **Shutdown Preference** in the PC Properties.

Base Snapshot	NONE (do not revert to snapshot)
Shutdown Preference	Graceful Shutdown from Operating System

- Update **Base Snapshot** to reflect the settings in the table below.

Virtual Machine	Base Snapshot
Firewall	NONE
Desktop	NONE


- Change **Shutdown Preference** to reflect the settings from the table below.

Virtual Machine	Shutdown Preference
Firewall	Keep Running
Desktop	Graceful Shutdown from Operating System

It is important to set the Gateway Pod Firewall to **Keep Running**, since it is acting as the gateway to the Internet for all FE pods.

It is also recommended that you leverage VMware vSphere’s virtual machine startup/shutdown to ensure that the gateway Firewall powers on when your host boots.

- Review the information on the screen and click **Update PC Settings**.




- Click on **Show Pod**.
- Repeat Steps 2-9 for the remaining virtual machines in the pod.

4.2.6 Bring the GW Master Pod Online

Follow the steps below to bring the GW Master pod online.

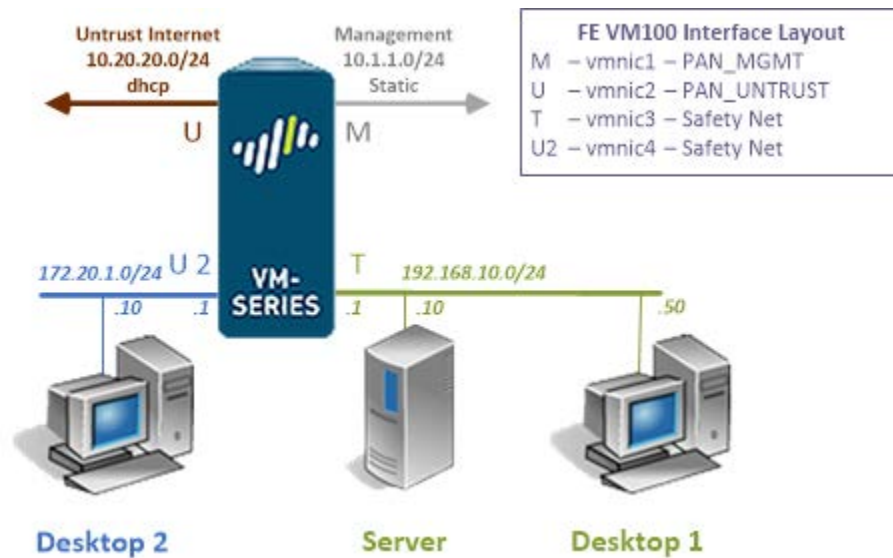
- Click on the **Equipment Pods** icon in the administrative interface.
- Click on the **Pod ID** hyperlink for your GW Master pod.
- Click on the **Online** button.



4.3 Firewall Essentials Master (FE) Pod Setup

The instructions in the sub-sections below will guide you through the installation of the Firewall Essentials (FE) pod. This pod requires the implementation and configuration of the Gateway pod prior to attempting licensing the Firewalls in the FE master and all clones.

The following topology shows two important links, first the untrust or U link is bound to the PAN_UNTRUST vSwitch and the management or M link is bound to the PAN_MGMT vSwitch. The other vSwitches are dynamically built by NETLAB+.




The PAN_MGMT and PAN_UNTRUST connections are essential since they are responsible for initial pod configuration, licensing and communication to the Internet.

4.3.1 Deploying FE Virtual Machine OVF/OVA Files

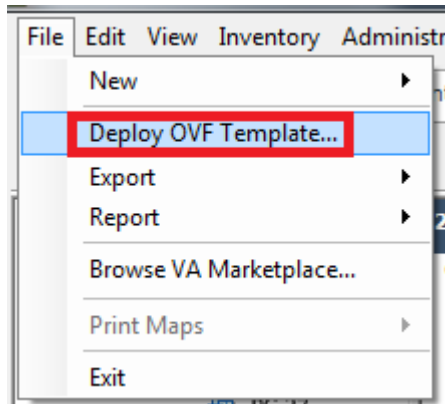
The OVF/OVA files that you downloaded from CSSIA must be deployed to your host server.

1. Open the **vClient** on your administration machine, ensure that your downloaded OVA/OVF files are accessible on this machine and then connect to your **vCenter Server**.
2. Select **Hosts and Clusters** in the address bar.



3. Click on the target ESXi Host Server.

- Click on **File -> Deploy OVF Template**.



- Click on **Browse** and locate the **PAN7_FE_FM_<VM Name>** OVF/OVA files you downloaded from CSSIA. Click **Next** to continue.
- On the OVF Template Details window, click **Next**.
- On the Name and Location window, change the name to **PAN7_FE_Master**. Click **Next**.
- On the Datastore window, select the target datastore and click **Next**.
- On the Disk Format window, select **Thin provisioned format** and click **Next**.
- Perform the following on the *Firewall virtual machine only*. Leave the default networks on non-Firewall virtual machines.
 - On the Firewall Essentials Firewall **Network Mapping** window, ensure the following networks are set and then click **Next**.

Virtual NIC	Setting
vmnic1	PAN_MGMT
vmnic2	PAN_UNTRUST
vmnic3	SAFETY NET
vmnic4	SAFETY NET


- On the **Ready to Complete** window, confirm the information and then click **Finish**.
- vCenter will begin deploying the virtual machine. This may take some time depending on the speed of your connection, HDDs, etc. When completed, click on **Close**.
- Perform the previous steps for each virtual machine in the FE pod.

4.3.2 Create Snapshots on the Master Virtual Machines

In order to proceed with pod cloning, snapshots must be created for the Master virtual machines.

Verify that all VMs are powered off before taking snapshots.

1. Open the **vClient** on your management workstation. Connect to your **vCenter Server**.
2. Select **Hosts and Clusters** in the address bar.



3. Right-click on each virtual machine and select **Snapshot > Take Snapshot**.
4. Enter **GOLDEN_MASTER** as the Snapshot Name.
5. Enter a description. It is a good idea to include the date in the description for later reference.
6. Click **OK**.
7. Repeat Steps 3-6 for the remaining virtual machines in the pod.
8. When all tasks have completed, log out of the vClient software.

4.3.3 NETLAB+ Virtual Machine Inventory Setup

This section will guide you in adding your templates to the Virtual Machine Inventory of your NETLAB+ system.

1. Login into your NETLAB+ system using the administrator account.
2. Select the **Virtual Machine Infrastructure** link.



[Virtual Machine Infrastructure](#)


3. Click the **Virtual Machine Inventory** link.




Virtual Machine Inventory

Import, clone, and manage the inventory of virtual machines to be used with NETLAB+.


4. Click the **Import Virtual Machines** button.



5. Select the checkbox next to your Palo Alto Networks Firewall Essentials virtual machines and click **Import Selected Virtual Machines**.



6. When the Configure Virtual Machines window loads, you can set your virtual machine parameters.
 - a. Check the drop-down box for the correct operating system for each imported virtual machine.
 - b. Select **Master** for the role of each virtual machine.
 - c. Add any comments for each virtual machine in the box to the right.
 - d. Verify your settings and click **Import Selected Virtual Machines**.



- e. Click OK when the virtual machines have finished loading.
- f. Verify that your virtual machines show up in the inventory.


4.3.4 Install the Master FE pod

This section will assist you in adding the Palo Alto Networks Firewall Essentials pod to your NETLAB+ system.


1. Login into NETLAB+ with the administrator account.
2. Select the **Equipment Pods** link.



3. Create a new pod by scrolling to the bottom and clicking the **Add a Pod** button.



4. On the New Pod Wizard page, click **Next**.
5. Then select the Palo Alto Networks Firewall Essentials pod radio button and click **Next**.



6. Select a Pod ID and click **Next**.

It is best practice to use a block of sequential ID numbers to number the Palo Alto Networks Firewall Essentials pods you are going to install. The Pod ID number determines the order in which the Palo Alto Networks Firewall Essentials pods will appear in the scheduler.


7. Type in **PAN7_FE_Master** for the Pod Name and click **Next**.
8. To finalize the wizard, click **OK**.

4.3.5 Update the Master Pod


1. In the Equipment Pods list, click on the Pod ID hyperlink to open the Pod Management interface.
2. Click on the Magnifying Glass icon next to the Desktop1 PC. Please note that your PC IDs may not match the graphic below.

POD 1001 - PCs AND SERVERS (click the GO buttons to reconfigure)						
GO	NAME	PC ID	STATUS	TYPE / VM	OPERATING SYSTEM	
	Desktop1	4102	ONLINE	PAN7_FE_Master_Desktop1	Linux	
	Firewall	4103	ONLINE	ABSENT		
	Server	4104	ONLINE	ABSENT		
	Desktop2	4105	ONLINE	ABSENT		

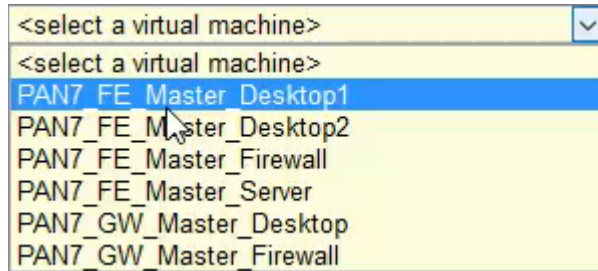
3. Click on **Modify PC Settings**.



4. Change the PC Type drop-down box to **Use Virtual Machine Inventory**.




- In the Base Virtual Machine window, select your Palo Alto Networks Firewall Essentials virtual machine to associate with the spot held for it in the pod.



- Update **Base Snapshot** to your *GOLDEN_MASTER* snapshot.
- Change **Shutdown Preference** to reflect the setting from the table below.

Virtual Machine	Shutdown Preference
Desktop1	Graceful Shutdown from Operating System
Desktop2	Graceful Shutdown from Operating System
Server	Graceful Shutdown from Operating System
Firewall	Graceful Shutdown from Operating System

- Review the information on the screen and click **Update PC Settings**.



- Click on **Show Pod**.
- Repeat Steps 2-9 for the remaining virtual machines in the pod.

Make sure the pod status is **Offline** prior to continuing. The cloning process requires the pod be offline. Since this is our master pod used for cloning other pods, we will keep it offline in order to use it to create the instructor and student pods.

5 Pod Cloning and Configuration

5.1 Pod Cloning

5.1.1 Linked Clones and Full Clones

NETLAB+ can create *linked clones* or *full clones*.

A **linked clone** (or linked virtual machine) is a virtual machine that shares virtual disks with the parent (or master) virtual machine in an ongoing manner. This conserves disk space and allows multiple virtual machines to common data. Linked clones can be created very quickly because most of the disk is shared with the parent VM.

A **full clone** is an independent copy of a virtual machine that shares nothing with the parent virtual machine after the cloning operation. Ongoing operation of a full clone is entirely separate from the parent virtual machine.

For the PAN7 FE Pods, we will create *linked* clone type with *normal* as the clone type.

The virtual machines have been preset with various configurations used to adjust the starting point for a particular lab. This allows us to reuse (revert to snapshot) the same virtual machines for various labs.

5.1.2 Creating User Pods

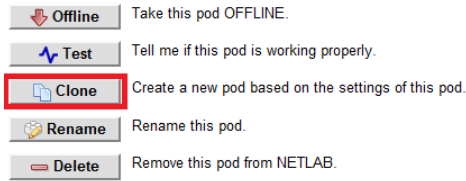
The following section describes how to create student pods on the same ESXi Host server as your Master Pod virtual machines. In this setup, we will create linked virtual machines using the NETLAB+ pod cloning utility.

1. Login into NETLAB+ with the administrator account.
2. Select the **Equipment Pods** link.



3. Click on your FE Master Pod.

- Click the **Clone** button to create a new pod based on the settings of this pod.



- Select the New Pod ID. It is advised to keep the pods in numerical order. If the pod IDs are not in numerical order, they will not show up in the scheduler in numerical order.
- Click **Next**.
- Enter a New Pod Name. For example, PAN7_FE_Pod1xxx. Click **Next**.

Using a structured naming convention for your pods will allow a better use of the PowerCLI scripts later in the install guide. The use of PowerCLI can shave hours off the installation time.

A good naming convention identifies the LabSet_PodVariance_Datastore_Host_PodID. Depending on your environment, you can trim this down to something like PAN7_FE_DS1_H2_P1021.

By doing this, you can programmatically identify the pods or virtual machines you plan manipulate by using wildcards. For instance, PAN7_FE_DS1_H2_P102* will select all PAN7 FE pods in the 102x range.

- When the action has finished processing, you are presented with a settings screen.

PC Name	Source Virtual Machine	Source Snapshot	→	Clone Name	Clone Type	Clone Role	Runtime Host or Group	Clone Datastore
Desktop1	PAN7_FE_Master_Desktop1	GOLDEN_MASTER		PAN7_FE_P1002 Desktop1	Linked	Normal	Host 192.168.37.22	local22
Firewall	PAN7_FE_Master_Firewall	GOLDEN_MASTER		PAN7_FE_P1002 Firewall	Linked	Normal	Host 192.168.37.22	local22
Server	PAN7_FE_Master_Server	GOLDEN_MASTER		PAN7_FE_P1002 Server	Linked	Normal	Host 192.168.37.22	local22
Desktop2	PAN7_FE_Master_Desktop2	GOLDEN_MASTER		PAN7_FE_P1002 Desktop2	Linked	Normal	Host 192.168.37.22	local22

- The three key columns for this Master Pod clone are Source Snapshot, Clone Type and Clone Role. The following settings should be applied to all 3 virtual machines:
 - Source Snapshot should be set to the **GOLDEN_MASTER** snapshot you created previously.
 - Under Clone Type, click the drop-down menu and verify that **Linked** is selected.
 - Under Clone Role, click the drop-down menu and select **Normal**.

10. When you are done changing settings, click **Clone Pod**. This should complete within a minute as we are creating linked virtual machines.
11. When the pod clone process is finished, click **OK**.

Time Saver: If you clone the 1st user pod instead of the Master pod, the defaults will all be set correctly and you will not have to change the Clone Type and Clone Role each time. NETLAB+ will still assume you want to link to the Master VMs, since Masters are ranked higher than Normal or Persistent VMs in the default pod cloning selections.


5.2 GW Pod Configuration

5.2.1 IP Address Assignment

There are two methods for applying an IP address to GW Firewall; Static or DHCP. We recommend that you use a static IP if possible to eliminate unknown and unexpected environment changes that may cause problems or even an outage to happen.

Many labs will require consistent Internet access to be completed and thus this pod needs to be treated as if was a service just like your NETLAB+ appliance.

To complete the instructions in the following sections, perform the following and then continue with the section that matches your configuration type.

1. Reserve the GW Master Pod.
2. Login to the GW Desktop machine using the username `sysadmin` and password of `Training$`.
3. Open a LX Terminal client window by clicking the  icon .
4. Press **Ctrl+Shift+T** to create a new tab the LX Terminal window. We will use one window to execute a local python program, which we will refer to as "*python terminal*" and the other to check the configuration of the Firewall as "*Firewall terminal*".

Having multiple tabs for python and Firewall is helpful in completing future steps.

5. In the Firewall terminal, log into the Firewall by typing the following command and then supplying the password `pa1oal1to` when prompted.

```
ssh admin@192.168.1.1
```

You will only need to perform one of the following sub-section instructions, either Static or DHCP, not both.

5.2.1.1 Static IP Address

Using a static IP address, as shown in this section, is the recommended method of applying an IP address to GW Firewall.

To provision a static IP address to your Firewall you will need to know the IP address, subnet mask and gateway IP so it can be assigned to ethernet1/1, the interface associated with vmnic 2 on the Firewall virtual machine.

IP Address _____
Subnet Mask _____
Gateway IP _____

Once you have these items, proceed with the following instruction set.

1. In your python terminal type the following, substituting your **IP Address**, **Subnet Mask** and **Gateway IP** respectively.

```
python3 panconfig.py gw static <IP Address> <Subnet Mask> <Gateway IP>
```

```
sysadmin@ubuntu:~$ python3 panconfig.py gw static 192.168.37.99 255.255.255.0 192.168.37.2
```

Only IPv4 addresses are supported for this interface.

2. In the Firewall terminal, type the following to display the interface information.

```
show interface all
```

3. Look specifically at the line with ethernet1/1 to verify that this matches the IP information that you had input. You may need to wait up to a minute for the system to commit and then display the proper change. Just press the up arrow key and then Enter to re-run your last command.
4. Now, we can perform a small test to ensure that the Firewall has default route outbound access via ethernet1/1, using a known IP address accessible via ICMP.

```
ping host 8.8.8.8
```

5.2.1.2 DHCP IP Address

The use of a dynamic IP address, as shown in this section, is an alternative method of applying an IP address to the GW Firewall. Skip this section if you have followed the recommended method using a static IP address, as described in the previous section.

The following procedure outlines the steps needed to allocate a dynamic IP address to your GW Firewall.

When using DHCP to provision your GW Firewall, we recommend that you create an IP reservation for the IP address in your DHCP pool that the Firewall will obtain.

1. In your python terminal, type the following command.

```
python3 panconfig.py gw dhcp
```

2. In the Firewall terminal, type the following to display the interface information.

```
show interface all
```

3. Look specifically at the line with ethernet1/1 to verify that this matches the IP information that you expect from your DHCP server. You may need to wait up to a minute for the system to commit and then display the proper change. Just press the up arrow key and then enter to re-run your last command.
4. Now, we can perform a small test to ensure that the Firewall has default route outbound access via ethernet1/1, using a known IP address accessible via ICMP.

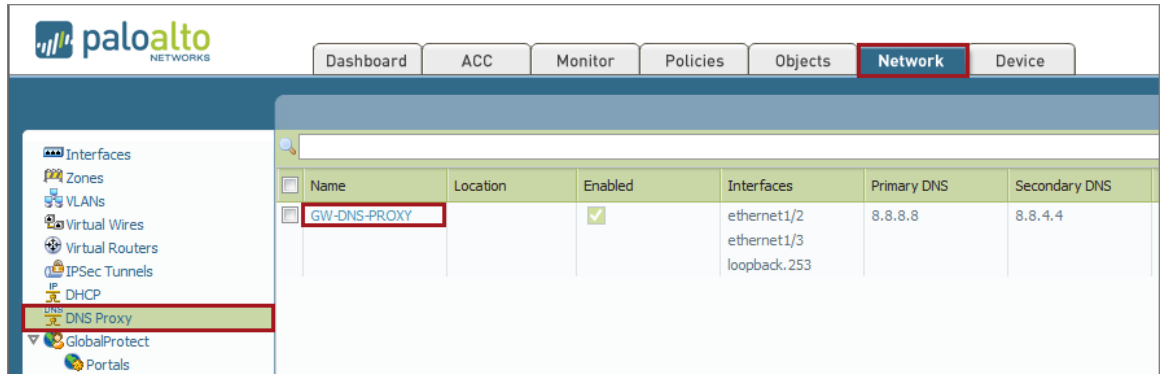
```
ping host 8.8.8.8
```

5.2.2 DNS Settings

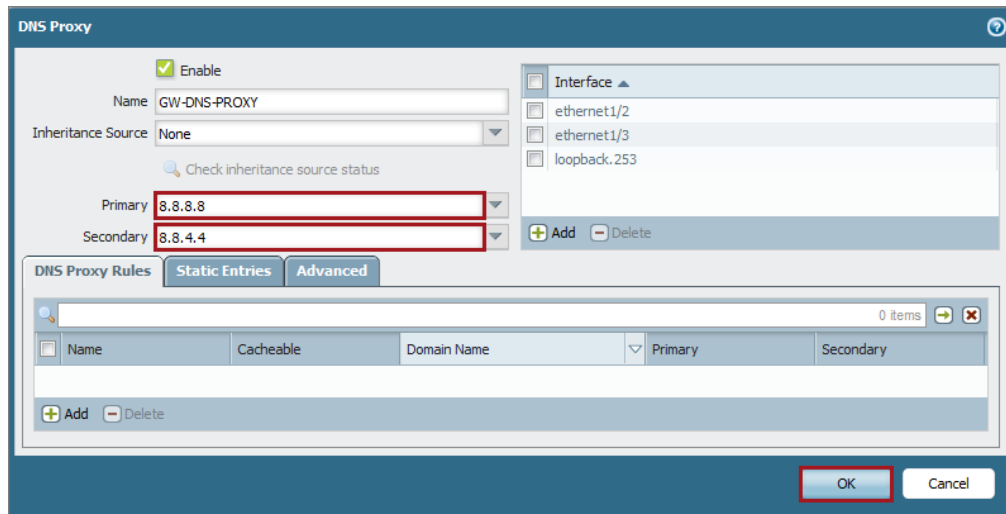
If your organization does not allow access to Google Public DNS, you will need to follow the procedure shown below to change the DNS to servers on the GW Firewall to those that you can access. The Google Public DNS server's IPv4 addresses are 8.8.8.8 and 8.8.4.4 respectively.

1. On the Gateway Desktop machine, open a new browser window, type `https://192.168.1.1` for the address and then press **Enter**.
2. Log in to the firewall using the following credentials; username `admin` and password `pa1oal1o`.

3. In the GW Firewall interface under the **Network > DNS Proxy**, click on the **GW-DNS-PROXY**.



4. In the DNS Proxy window, change your **Primary** and **Secondary** DNS settings as necessary and then click OK.



5. Once your settings have been set, **Commit** your changes and continue to the next section.



You can let the pod reservation end on its own since the firewall will remain running.

5.2.3 Licensing

Licenses can be obtained by following the instructions in Section 3.2.

Once you have your licenses, you are ready to activate licenses on the Gateway Pod.

1. In the python terminal, type the following command to license your GW Firewall. Remember to replace the Auth ID with the *auth id* you received from Palo Alto Network Academy.

```
python3 panconfig.py gw license <Auth ID>
```

When using the panconfig.py program you can use the `-h` switch to get help with the command. Try `panconfig.py -h` or `panconfig.py gw -h` to learn more about the commands.

2. When the license is pulled from Palo Alto Networks licensing service, the Firewall will disconnect all terminal sessions as it automatically reboots. This process will take between 3 and 4 minutes to complete.
3. Log back into the Firewall terminal and then type the following command into the CLI.

```
show system info
```

4. In the show system info output, look for an item name serial and verify that you have a 12-digit number; if so, your Firewall is licensed.

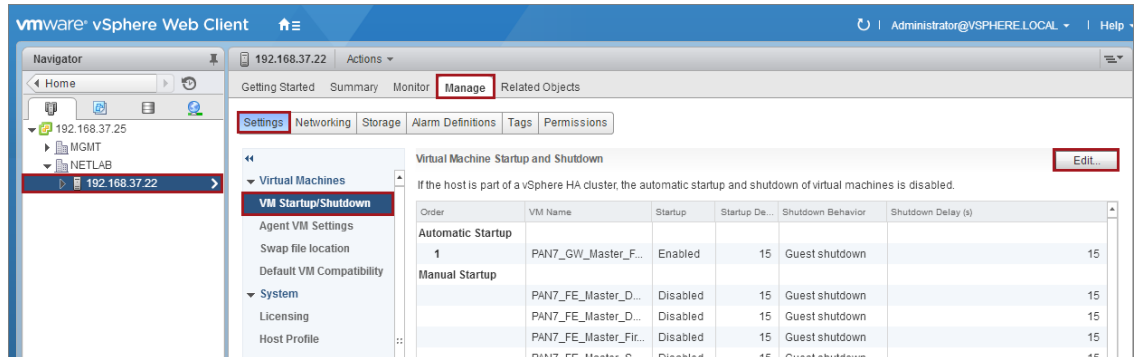
5.2.4 Startup and Shutdown the Firewall

In the following procedure, we use VMware ESXi to automate the startup and shutdown of the GW Firewall with the power cycle of the ESXi host to ensure that the Firewall service is brought up with the host.

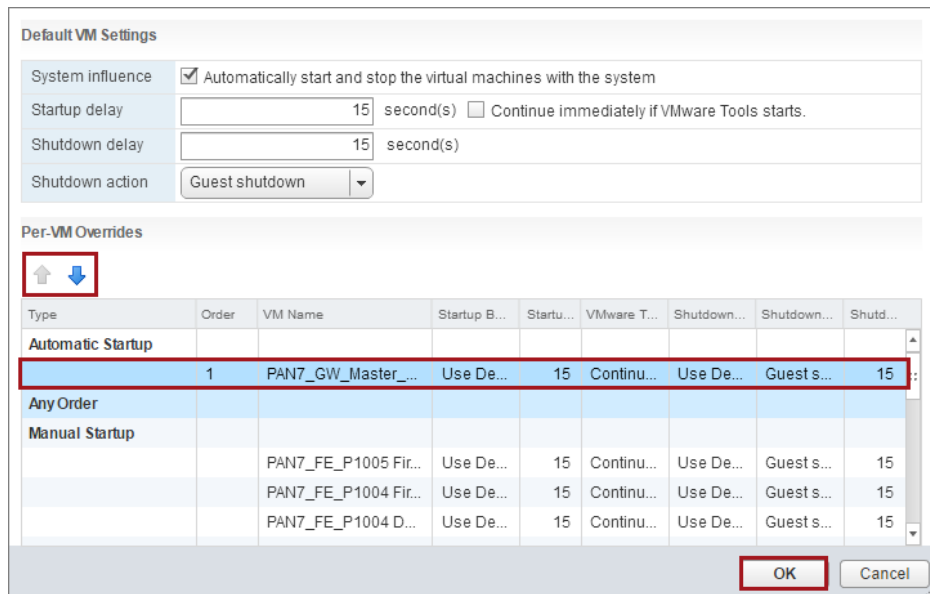
Remember, the Gateway Firewall is providing a service that is required for all FE pods that are running on that ESXi host. Without the GW Firewall up and running, the tethered FE pods will not have Internet access, which in turn will cause problems executing labs or lab functionality.

1. Log into your vCenter server.
2. In vCenter, click on **Hosts and Clusters**.

3. Select the **host** of your target GW Firewall virtual machine is on and then click **Manage, Settings, VM Startup/Shutdown, and Edit...**



4. Here, you can select your **PAN7_GW_Master** virtual machine and use the arrow icons to move it to the **Automatic Startup** section and then click **OK**.



5.3 FE Pod Configuration

The following instructions will have two distinctly different methods to complete the work that needs to be done prior to the cloned FE pods actual use.

For expediency as well as consistency, we have provided methods that leverage the vSphere PowerCLI. These methods will be identified along with the manual methods for accomplishing the same tasks. **Please do not execute both sets of tasks, since it would cause problems booting pods.**

The FE Master Pod does not need a license if it is only going to be used as a master and not intended for access by instructor or learner.

5.3.1 IP Addressing

IP addresses for the FE Firewalls are obtained by the GW Firewall via DHCP. The following sub-sections address the power-on requirements of the FE Firewalls to ensure proper DHCP fulfillment and ways that you can test the environment.


A 10-second or greater pause between boot of each FE Firewall will allow enough time between discovery requests to eliminate offer hijacking.

5.3.1.1 Boot FE Firewalls - Manual Method

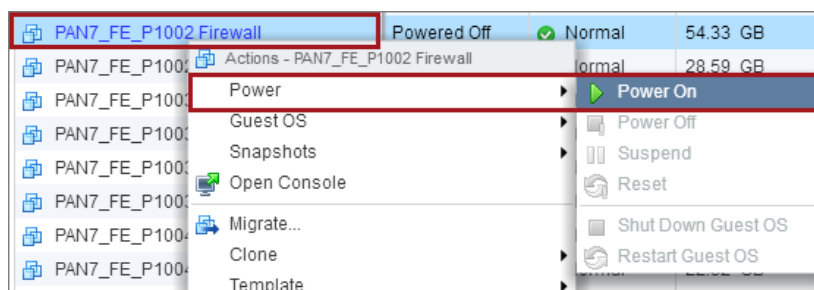
Choose one method; do not perform both methods. You may boot the FE Firewalls using the manual method described in this section or by using PowerCLI as described in the next section.

Manual boot will require that you access your ESXi host via the Web Client or vSphere Client and perform a manual boot in >10 second intervals. The following process will step you through booting FE Firewalls via the Web Client.

1. Access your vCenter administrative interface, either the vSphere Client or the Web Client.
2. Click on Hosts and Clusters then select the host your FE virtual machines reside on and click on **Related Objects** then click on **Virtual Machines**.



3. Now, right-click on the target Firewall and select **Power > Power On**.



4. Wait for at least 10 seconds before executing the previous step on the next of the Firewalls in the list. Continue until all cloned Firewalls (not the masters) are up and running. Then wait 5 minutes before continuing to Licensing.

5.3.1.2 Boot FE Firewalls - PowerCLI Method

Choose one method; do not perform both methods. You may boot the FE Firewalls using PowerCLI as described in this section or by using the manual method described in the previous section.

The following identifies how to perform a 10-second incremented boot of the FE Firewall virtual machines using PowerCLI.

1. Click on vSphere PowerCLI link to open a PowerCLI terminal window.
2. Login to your vCenter using the connection string below, replacing the appropriate placeholders with ones that match your configuration.

```
Connect-VIServer -Server <VCSA IP Address> -Protocol https -User
<username> -Password <password>
```

3. Once you are connected to your vCenter Server, you can use the following program to boot all of your FE Firewalls at an increment of 1 every 10 seconds.

```
$vms = get-vm "PAN7_FE_P*Firewall"
foreach($vm in $vms)
{
echo $vm
sleep 10
Start-VM $vm -Confirm:$false
}
```

Notice the name of the virtual machine in the first line, "PAN7_FE_P*Firewall". You will want to change this to match the naming convention you had used in naming your cloned pods.

4. Once the program has finished and all FE Firewalls are booting, wait for 5 minutes before proceed to Licensing.

5.3.2 Licensing

To complete the licensing of the Firewalls, proceed with the following steps

This procedure for licensing was created as a shortcut to the alternative manual configuration. You will need to be patient, since executing this program will take time to work through all the Firewalls that you have cloned out, especially if there are problems or the Firewall has not completed the boot process.

1. First, we will check to see if all Firewalls are responsive and not licensed by executing the following command in the GW Desktop python terminal.

```
python3 panconfig.py fe license_info -D
```

```
sysadmin@lubuntu:~$ python3 panconfig.py fe license_info -D
```

2. Monitor the output of the command. Normally, if you see that the IP Address: x.x.x.x not accessible like that which is listed below, then not all of the FE Firewalls are responding yet. This will require another run in a few minutes.

```
sysadmin@lubuntu:~$ python3 panconfig.py fe license_info -D
IP Address: 10.20.20.1 not accessible. Please check connection.
IP Address: 10.20.20.2 not accessible. Please check connection.
Executed:10.20.20.3      Serial:unknown
Executed:10.20.20.4      Serial:unknown
Executed:10.20.20.5      Serial:unknown
End of DHCP list.
```

3. You will need to track the count of the FE Firewalls that you have created and powered on. In this demonstrated scenario, we have 4 new Firewalls, the first IP address was residual in the GW Firewall DHCP server from previous use. Our complete list then looked like this:

```
sysadmin@lubuntu:~$ python3 panconfig.py fe license_info -D
IP Address: 10.20.20.1 not accessible. Please check connection.
Executed:10.20.20.2      Serial:unknown
Executed:10.20.20.3      Serial:unknown
Executed:10.20.20.4      Serial:unknown
Executed:10.20.20.5      Serial:unknown
End of DHCP list.
```

You can use the panconfig tool to look at the DHCP leases by typing the following command: `python3 panconfig.py gw dhcp_lease`

You can also access and display the DHCP list by executing the following in the Firewall terminal: `show dhcp server lease ethernet1/2`

4. After confirming that all Firewalls are up and responding, run the following command to license all Firewalls using a single auth id.

```
python3 panconfig.py fe license -D <Auth ID>
```

Each Firewall will perform a reboot after receiving a license from the Palo Alto service. The reboot process takes about 4 minutes to complete, so please allow enough time after the last system is licensed before proceeding to the next step.

5. The output should look similar the following screenshot.

```
sysadmin@ubuntu:~$ python3 panconfig.py fe license -D [REDACTED]
IP Address: 10.20.20.1 not accessible. Please check connection.
Processing auth code [REDACTED] on 10.20.20.2
Processing auth code [REDACTED] on 10.20.20.3
Processing auth code [REDACTED] on 10.20.20.4
Processing auth code [REDACTED] on 10.20.20.5
End of DHCP list.
```

6. Confirm that all Firewalls have been licensed by executing the following command.

```
python3 panconfig.py fe license_info -D
```

7. The output will look similar to the screenshot below when responding correctly. Notice that the output now displays actual serial numbers for all of the Firewalls queried.

```
sysadmin@ubuntu:~$ python3 panconfig.py fe license_info -D
IP Address: 10.20.20.1 not accessible. Please check connection.
Executed:10.20.20.2      Serial:007000010929
Executed:10.20.20.3      Serial:007000010930
Executed:10.20.20.4      Serial:007000010931
Executed:10.20.20.5      Serial:007000010932
End of DHCP list.
```

If you have successfully completed the previous steps, proceed to the Pod Snapshots Section.

If you have not successfully completed the previous steps, refer to the next section, **Troubleshooting**.

5.3.2.1 Troubleshooting

There are various tools at our disposal to help find answers to the problems that we may encounter when setting up these pods. It may be helpful to review the objectives of this installation as a refresher.

- The previous steps require Internet access to be successful.

You can check Internet access by opening a web browser on the GW Desktop machine and surfing to Google or Yahoo. The default route for the Desktop machine is the PAN_UNTRUST, which the FE Firewalls use to get out to the Internet, making it a good place to test.

- The FE Firewalls must be able to obtain an IP address from the GW Firewall. The communication required for this to happen needs to be on specific interfaces and the port groups (if setup on the same vswitch) must have vlan ids assigned.

Interface mapping:

```

PAN7 GW Firewall:
  vmnic1 -> management -> SAFETY NET -> NDG vSwitch
  vmnic2 -> ethernet1/1 -> VM Network -> DMZ/Uplink for Internet
Access
  vmnic3 -> ethernet1/2 -> PAN_UNTRUST -> PAN FE Untrust
  vmnic4 -> ethernet1/3 -> PAN_MGMT -> PAN FE Management

PAN7 GW Desktop:
  vmnic1 -> eth0 -> SAFETY NET -> NDG vSwitch
  vmnic2 -> eth1 -> PAN_UNTRUST -> PAN FE Untrust

PAN7 FE Firewall:
  vmnic1 -> management -> PAN_MGMT -> PAN FE Management
  vmnic2 -> ethernet1/1 -> PAN_UNTRUST -> PAN FE Untrust
  vmnic3 -> ethernet1/2 -> SAFETY NET -> PAN FE Trust
  vmnic4 -> ethernet1/3 -> SAFETY NET -> PAN FE Untrust2 (DMZ??)

```

- The FE Firewall vmnic2 and GW Firewall vmnic3 need to be set to the PAN_UNTRUST network.
- If vlan ids are not assigned, another interface that isn't used by default will request an IP address causing dual IPs per FE Firewall.

5.3.3 Pod Snapshots

Before we create snapshots, we will need to shutdown each FE Firewall that we have running. It is preferred to perform a “graceful” shutdown of each system. Once the systems are shutdown, we can perform snapshots on all systems.

We left the snapshots to the very last portion to make it easy for those who wished to use PowerCLI to speed up the process.


Gracefully shutdown the FE Firewalls by performing the following steps. After that, you will create pod snapshots using one of the two methods discussed in the sub-sections below.

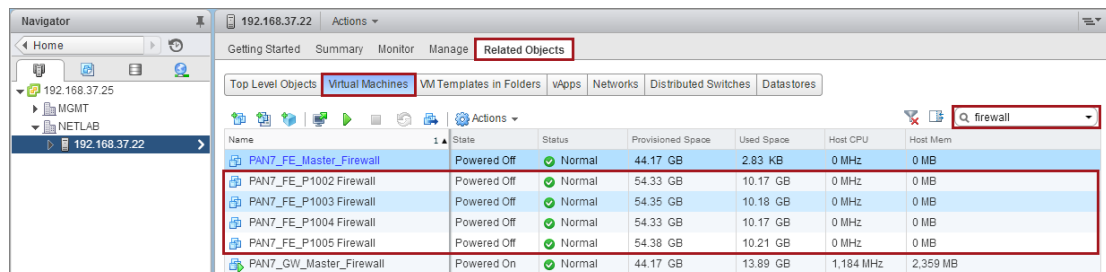
1. In the GW Desktop python terminal type in the following command to shut the FE Firewall down.

```
python3 panconfig.py fe shutdown -D
```

2. Observe the output of the program to verify the program communicated with each Firewall and initiated a shutdown command on each. The output should resemble that of the screenshot below.

```
sysadmin@lubuntu:~$ python3 panconfig.py fe shutdown -D
IP Address: 10.20.20.1 not accessible. Please check connection.
System 10.20.20.2 is being shutdown.
System 10.20.20.3 is being shutdown.
System 10.20.20.4 is being shutdown.
System 10.20.20.5 is being shutdown.
End of DHCP list.
```

3. Observe the status of the virtual machines in vCenter via the vSphere Client or Web Client to ensure they are no longer powered on. You may need to click the refresh button  as the web client seldom refreshes autonomously. Also, you can specify a name in the filter section to narrow the list of virtual machines displayed. Here we are only focused on the Firewall’s so we put Firewall in the filter text box.



Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
PAN7_FE_Master_Firewall	Powered Off	Normal	44.17 GB	2.83 KB	0 MHz	0 MB
PAN7_FE_P1002 Firewall	Powered Off	Normal	54.33 GB	10.17 GB	0 MHz	0 MB
PAN7_FE_P1003 Firewall	Powered Off	Normal	54.35 GB	10.18 GB	0 MHz	0 MB
PAN7_FE_P1004 Firewall	Powered Off	Normal	54.33 GB	10.17 GB	0 MHz	0 MB
PAN7_FE_P1005 Firewall	Powered Off	Normal	54.38 GB	10.21 GB	0 MHz	0 MB
PAN7_GW_Master_Firewall	Powered On	Normal	44.17 GB	13.89 GB	1,184 MHz	2,359 MB

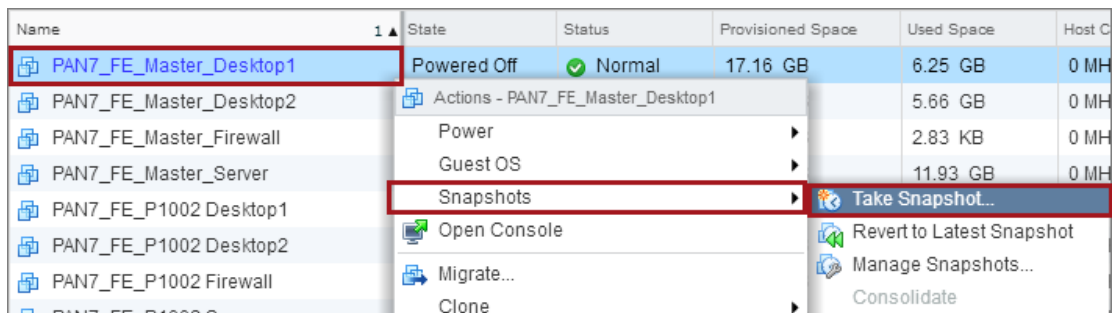
Once all FE Firewall virtual machines are in a “Powered Off” state you may snapshot the virtual machines using one of the two methods described below.

5.3.3.1 Snapshot the Virtual Machines - Manual Method

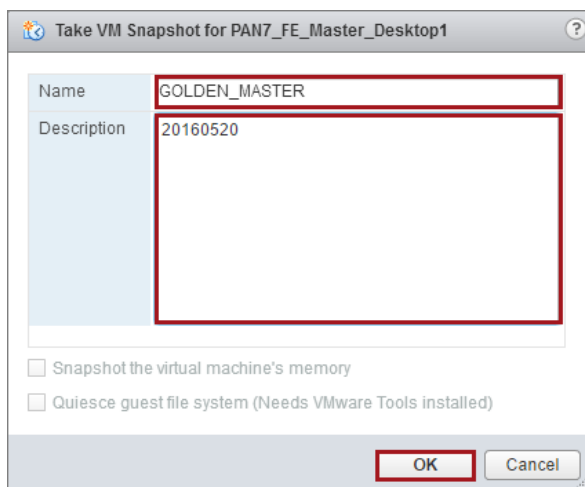
Choose one method: This section provides details on manually creating snapshots of your virtual machines. You may use this method or follow the instructions in the next section to create snapshots using PowerCLI.

Manually snapshotting a virtual machine can be accomplished in a variety of ways including vSphere Client, vSphere Web Client and NETLAB+. The following procedure highlights the use of the vSphere Web Client method, since it is likely the one in most use on newer installs.

1. Using the vSphere Web Client, right-click on the target virtual machine and select **Snapshots -> Take Snapshot**.



2. In the **Take VM Snapshot** window, type **GOLDEN_MASTER** in the **Name** text box and type today's date in the **Description** and then click **OK**.



3. Repeat the previous steps for all virtual machines in the cloned FE pods.

5.3.3.2 Snapshot the Virtual Machines - PowerCLI Method

Choose one method: This section provides details on creating snapshots of your virtual machines using PowerCLI. You may use this method or follow the instructions in the previous section to manually create snapshots.

Using PowerCLI can be tedious and require time to learn or understand, but if used well can save a large amount of time on your installation. The following procedure will leverage PowerCLI to create snapshots of all FE pod virtual machines.

The key for this procedure to work well is in the naming convention that was used for the FE pods. As long as the virtual machines can be selected accurately, using wildcards then this procedure will work without issue.

1. Click on vSphere PowerCLI link to open a PowerCLI terminal window.
2. Login to your vCenter using the connection string below, replacing the appropriate placeholders with ones that match your configuration.

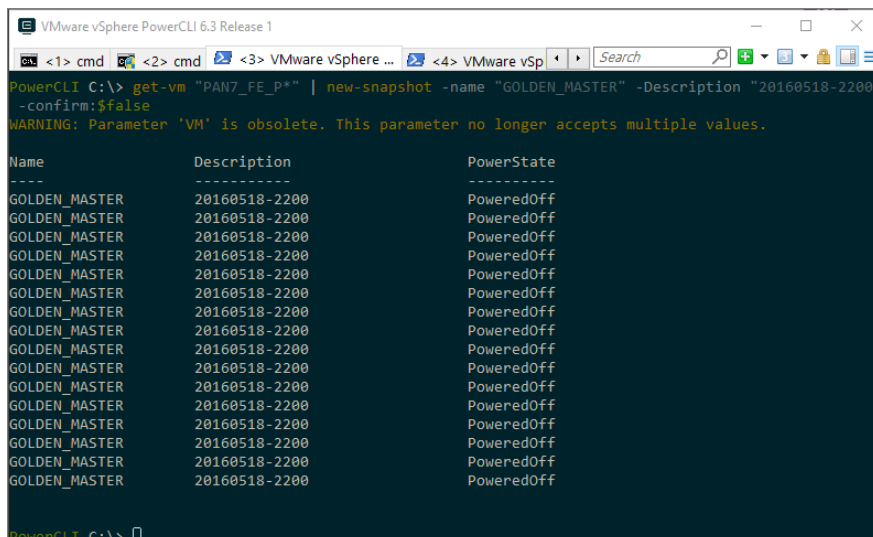
```
Connect-VIServer -Server <VCSA IP Address> -Protocol https -User  
<username> -Password <password>
```

3. Type the following command into your PowerCLI window, replacing the **<VM_Name>** and **<DateStamp>** with information pertinent to your installation.

Remember that if you used the naming convention recommended that you will need to use it with a wildcard (similar to **PAN7_FE_P***) to select and apply the new snapshot the returned virtual machine instances.

```
get-vm "<VM_Name>" | new-snapshot -name "GOLDEN_MASTER" -Description  
"<DateStamp>" -confirm:$false
```

4. Observe that your output looks similar to the following screenshot.




```
VMware vSphere PowerCLI 6.3 Release 1  
PowerCLI C:\> get-vm "PAN7_FE_P*" | new-snapshot -name "GOLDEN_MASTER" -Description "20160518-2200"  
-confirm:$false  
WARNING: Parameter 'VM' is obsolete. This parameter no longer accepts multiple values.  
Name Description PowerState  
----  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
GOLDEN_MASTER 20160518-2200 PoweredOff  
PowerCLI C:\>
```

5.4 Bring Pods Online

1. Login into NETLAB+ with the administrator account.
2. Select the **Equipment Pods** link.



3. Click on a cloned FE pod that was created.
4. Click the **Online** button bring this pod online for use.



5. Repeat the previous steps for each of the cloned FE pods that you wish to make available to your learners.

6 PAN Firewall Administration Best Practices

6.1 Administration

We recommend that you use the *PAN-OS Administrator's Guide v7* to assist you or your firewall administrator with the configuration and operation of the Gateway Firewall (VM100).

Listed below are some administration topics that you will likely want to consider adjusting in your system.

6.2 Security Policies

Traditional firewalls typically have a Deny All rule at the end of the firewall rules. Placing this rule on a Palo Alto Networks platform will break your traffic. The Palo Alto Networks platform has two policies at the end of the firewall, Intrazone and Interzone, which handle this same function. It is better to enable logging on the Interzone policy. If you become concerned with the traffic between pods, then you could enable logging on the Intrazone policy. Details on how to enable logging on these policies can be found in the Administrators Guide for the firewall.

Another best practice for security policy creation is to use Application Default for the services on policies that allow and Any for the services on policies that deny.

6.3 Logging

Logging is an essential component when providing Internet access. Logging provides the mechanism for recording Internet transactions, sessions, MAC address, user events, etc. Many of the items necessary for tracing any malicious activity from either untrusted or trusted locations.

It is recommended that you create a separate storage facility for logging, especially if you have installed this lab set on multiple hosts. Having a separate logging facility will allow you to aggregate logs from multiple systems, allow for longer log retention capabilities as well as providing various search features.

More information regarding the use and configuration PAN-OS v7 with a syslog server can be found in the *PAN-OS v7 Administration Guide* under a section titled, *Use Syslog for Monitoring* as well as the Palo Alto Networks Firewall Essentials Lab 15.

6.4 Threat Prevention

6.4.1 URL Filtering

Depending upon your organization's governance, policies and guidelines, you may need to increase the restrictions regarding access to and from your learner pods. URL Filtering can be used to block everything except URLs that your organization allows. More information regarding the use and configuration of URL Filtering can be found in the *PAN-OS v7 Administration Guide* under a section titled *URL Filtering*.

6.4.2 Wildfire

Wildfire can identify and protect your users from malware. It is important that you ensure this service is configured and running correctly on your Gateway Firewall. Additional documentation for Wildfire can be found in the *Wildfire™ Administrators Guide Version 7.0*.

6.4.3 Monitoring

It is important to monitor your platform on a regular basis. The best way to monitor the system is from the Monitor tab in the GUI. The logs to especially pay attention to is the Threats, Wildfire and URL filtering. In addition to those, you will want to check the Traffic log. With the Traffic log, the administrator needs to become comfortable with the regex filtering. Information on filtering can be found in the Administrator Guide. Some things to look for are Deny action in the logs. Pay close attention to the policies that caused the Deny action. The Wildfire logs should be monitored for Malware verdicts and URL filtering should be adjusted to match your school's policies.