

vMX

vMX Getting Started Guide for VMware

Published
2021-06-10

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

vMX vMX Getting Started Guide for VMware
Copyright © 2021 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | vi

1

vMX Overview

vMX Overview | 2

Virtual Network Interfaces for vMX | 7

2

Prepare vMX Installation on VMWare

Minimum Hardware and Software Requirements | 11

vMX Package Contents | 15

3

Installing vMX

Install vMX on VMWare | 18

Installing vMX with OVA Files | 18

Setting Up the Datastore | 19

Setting Up the Network | 20

Deploying the VCP VM | 24

Deploying the VFP VM | 25

Launching vMX on VMware | 28

Manage vMX on VMWare | 28

Allowing VLAN Traffic | 29

Setting Up a Serial Port Connection to the VM | 29

Viewing CPU Information | 32

Verifying Whether VMs Are Running | 32

Enabling SR-IOV for VMware | 33

Enable SR-IOV on a Physical Adapter | 33

Add vSwitch and Port Groups | 35

Assigning the SR-IOV NIC to the VFP VM | 36

Enable VLAN Tagging on SR-IOV Interfaces | 39

SR-IOV Interface for Layer 2 Services | 41

4

Configuring vMX Chassis-Level Features

Configuring the Number of Active Ports on vMX | 45

Naming the Interfaces | 45

Configuring the Media MTU | 46

Enabling Performance Mode or Lite Mode | 47

Tuning Performance Mode | 49

5

Class of Service for vMX

CoS on vMX Overview | 52

CoS Features and Limitations on vMX | 54

Configuring Hierarchical CoS on vMX | 56

Enabling Flexible Queuing | 56

Mapping Forwarding Classes to Queues on vMX | 57

Configuring Traffic Control Profiles for vMX | 57

Configuring Schedulers on vMX | 57

Example: Configuring Hierarchical CoS on vMX | 59

Requirements | 59

Overview | 59

Configuration | 59

Configuring Four-Level Hierarchical Scheduling on vMX | 65

Packet Loss Priority and Drop Profiles on vMX | 66

Managing Congestion Using Drop Profiles and Packet Loss Priorities on vMX | 68

Configuring Drop Profiles | 68

Configuring Schedulers with Drop Profiles | 69

6

Troubleshooting vMX

[Viewing VFP Statistics | 72](#)

[Viewing VFP Log Files | 73](#)

[Troubleshooting VFP and VCP Connection Establishment | 74](#)

About This Guide

Use this guide to install the virtual MX router in the VMware environment. This guide also includes basic vMX configuration and management procedures.

After completing the installation and basic configuration procedures covered in this guide, refer to the Junos OS documentation for information about further software configuration on the vMX router.

RELATED DOCUMENTATION

| [Junos OS for MX Series Documentation Home](#)

1

CHAPTER

vMX Overview

[vMX Overview | 2](#)

[Virtual Network Interfaces for vMX | 7](#)

vMX Overview

IN THIS SECTION

- [Benefits and Uses of vMX Routers | 2](#)
- [Automation for vMX Routers | 3](#)
- [Architecture of a vMX Instance | 3](#)
- [Traffic Flow in a vMX Router | 6](#)

Read this topic to get an overview about vMX virtual routers.

The vMX router is a virtual version of the MX Series 3D Universal Edge Router. Like the MX Series router, the vMX router runs the Junos operating system (Junos OS) and supports Junos OS packet handling and forwarding modeled after the Trio chipset. Configuration and management of vMX routers are the same as for physical MX Series routers, allowing you to add the vMX router to a network without having to update your operations support systems (OSS).

You install vMX software components on an industry-standard x86 server running a hypervisor, either the kernel-based virtual machine (KVM) hypervisor or the VMware ESXi hypervisor.

For servers running the KVM hypervisor, you also run the Linux operating system and applicable third-party software. vMX software components come in one software package that you install by running an orchestration script included with the package. The orchestration script uses a configuration file that you customize for your vMX deployment. You can install multiple vMX instances on one server.

For servers running the ESXi hypervisor, you run the applicable third-party software.

Some Junos OS software features require a license to activate the feature. To understand more about vMX Licenses, see, [vMX Licenses for KVM and VMware](#). Please refer to the [Licensing Guide](#) for general information about License Management. Please refer to the product [Data Sheets](#) for further details, or contact your Juniper Account Team or Juniper Partner.

Benefits and Uses of vMX Routers

You can use virtual devices to lower your capital expenditure and operating costs, sometimes through automating network operations. Even without automation, use of the vMX application on standard x86 servers enables you to:

- Quickly introduce new services
- More easily deliver customized and personalized services to customers
- Scale operations to push IP services closer to customers or to manage network growth when growth forecasts are low or uncertain
- Quickly expand service offerings into new sites

A well designed automation strategy decreases costs as well as increasing network efficiency. By automating network tasks with the vMX router, you can:

- Simplify network operations
- Quickly deploy new vMX instances
- Efficiently install a default Junos OS configuration on all or selected vMX instances
- Quickly reconfigure existing vMX routers

You can deploy the vMX router to meet some specific network edge requirements, such as:

- Network simulation
- Terminate broadband subscribers with a virtual broadband network gateway (vBNG)
- Temporary deployment until a physical MX Series router is available

Automation for vMX Routers

Automating network tasks simplifies network configuration, provisioning, and maintenance. Because the vMX software uses the same Junos OS software as MX Series routers and other Juniper Networks routing devices, vMX supports the same automation tools as Junos OS. In addition, you can use standard automation tools to deploy the vMX, as you do other virtualized software.

Architecture of a vMX Instance

The vMX architecture is organized in layers:

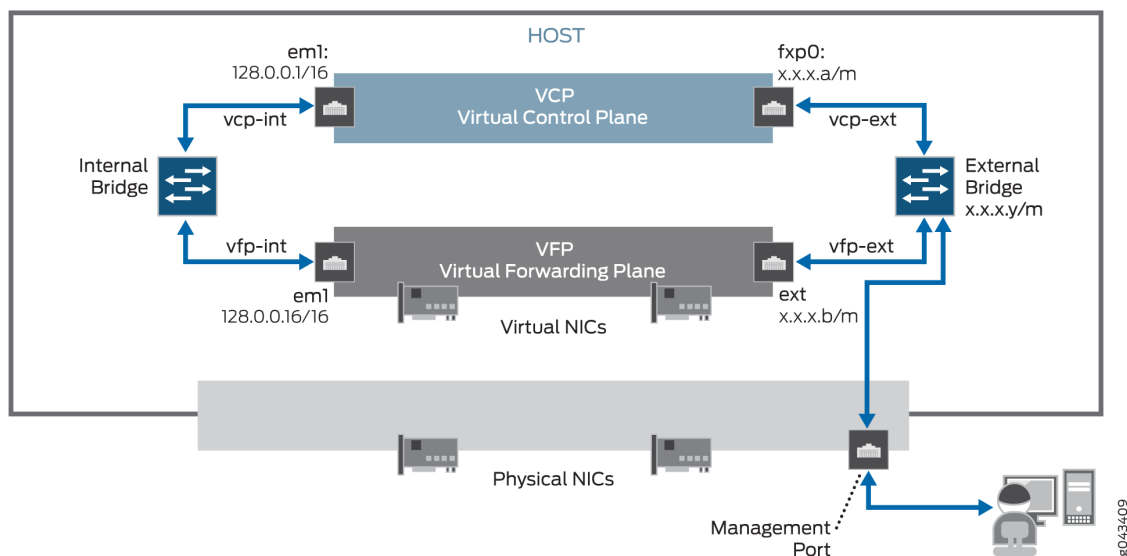
- The vMX router at the top layer
- Third-party software and the hypervisor in the middle layer

Linux, third-party software, and the KVM hypervisor in the middle layer in Junos OS Release 15.1F3 or earlier releases. In Junos OS Release 15.1F3 and earlier releases, the host contains the Linux operating system, applicable third-party software, and the hypervisor.

- The x86 server in the physical layer at the bottom

Figure 1 on page 4 illustrates the architecture of a single vMX instance inside a server. Understanding this architecture can help you plan your vMX configuration.

Figure 1: vMX Instance in a Server



The physical layer of the server contains the physical NICs, CPUs, memory, and Ethernet management port. The host contains applicable third-party software and the hypervisor.


Supported in Junos OS Release 15.1F3 and earlier releases, the host contains the Linux operating system, applicable third-party software, and the hypervisor.

The vMX instance contains two separate virtual machines (VMs), one for the virtual forwarding plane (VFP) and one for the virtual control plane (VCP). The VFP VM runs the virtual Trio forwarding plane software and the VCP VM runs Junos OS.

The hypervisor presents the physical NIC to the VFP VM as a virtual NIC. Each virtual NIC maps to a vMX interface. Figure 2 on page 5 illustrates the mapping.

The orchestration script maps each virtual NIC to a vMX interface that you specify in the configuration file. After you run the orchestration script and the vMX instance is created, you use the Junos OS CLI to configure these vMX interfaces in the VCP (supported in Junos OS Release 15.1F3 or earlier releases).

Figure 2: NIC Mapping



After the vMX instance is created, you use the Junos OS CLI to configure these vMX interfaces in the VCP. The vMX router supports the following types of interface names:

- Gigabit Ethernet (ge)
- 10-Gigabit Ethernet (xe)
- 100-Gigabit Ethernet (et)

NOTE: vMX interfaces configured with the Junos OS CLI and the underlying physical NIC on the server are independent of each other in terms of interface type (for example, ge-0/0/0 can get mapped to a 10-Gigabit NIC).

The VCP VM and VFP VM require Layer 2 connectivity to communicate with each other. An *internal* bridge that is local to the server for each vMX instance enables this communication.

The VCP VM and VFP VM also require Layer 2 connectivity to communicate with the Ethernet management port on the server. You must specify virtual Ethernet interfaces with unique IP addresses

and MAC addresses for both the VFP and VCP to set up an *external* bridge for a vMX instance. Ethernet management traffic for all vMX instances enters the server through the Ethernet management port.

The way network traffic passes from the physical NIC to the virtual NIC depends on the virtualization technique that you configure.

vMX can be configured to run in two modes depending on the use case:

- Lite mode—Needs fewer resources in terms of CPU and memory to run at lower bandwidth.
- Performance mode—Needs higher resources in terms of CPU and memory to run at higher bandwidth.

NOTE: Performance mode is the default mode.

Traffic Flow in a vMX Router

The x86 server architecture consists of multiple sockets and multiple cores within a socket. Each socket also has memory that is used to store packets during I/O transfers from the NIC to the host. To efficiently read packets from memory, guest applications and associated peripherals (such as the NIC) should reside within a single socket. A penalty is associated with spanning CPU sockets for memory accesses, which might result in non-deterministic performance.

The VFP consists of the following functional components:

- Receive thread (RX): RX moves packets from the NIC to the VFP. It performs preclassification to ensure host-bound packets receive priority.
- Worker thread: The Worker performs lookup and tasks associated with packet manipulation and processing. It is the equivalent of the lookup ASIC on the physical MX Series router.
- Transmit thread (TX): TX moves packets from the Worker to the physical NIC.

The RX and TX components are assigned to the same core (I/O core). If there are enough cores available for the VFP, the QoS scheduler can be allocated separate cores. If there are not enough cores available, the QoS scheduler shares the TX core.

TX has a QoS scheduler that can prioritize packets across several queues before they are sent to the NIC (supported in Junos OS Release 16.2).

The RX and TX components can be dedicated to a single core for each 1G or 10G port for the most efficient packet processing. High-bandwidth applications must use SR-IOV. The Worker component utilizes a scale-out distributed architecture that enables multiple Workers to process packets based on

packets-per-second processing needs. Each Worker requires a dedicated core (supported in Junos OS Release 16.2).

RELATED DOCUMENTATION

[Virtual Network Interfaces for vMX | 7](#)

https://www.juniper.net/documentation/en_US/release-independent/licensing/topics/topic-map/vmx-licensing.html

Virtual Network Interfaces for vMX

IN THIS SECTION

- [Paravirtualization | 8](#)
- [PCI Passthrough with SR-IOV | 9](#)

In a virtual environment, packet input and output capabilities play a significant role in the performance of the packet processing functionality inside the virtual machine, specifically the VFP VM. VFP supports two types of virtual network interfaces:

- **Paravirtualized**—Paravirtualized network interfaces use network drivers in the guest OS and host OS that interact with the virtual environment and communicate effectively to give higher performance than fully emulated interfaces. In KVM, the supported paravirtualized interface is virtio. For VMware, VMXNET3 is supported.
- **PCI pass-through**—PCI pass-through enables PCI devices such as network interfaces to appear as if they were physically attached to the guest operating system, bypassing the hypervisor and providing a high rate of data transfer. The physical network interfaces support single root I/O virtualization (SR-IOV) capability and can be connected to the VMs using PCI pass-through.

Choose the type based on how you want to use the vMX router. See [Table 1 on page 8](#).

Table 1: Considerations for Choosing a Virtualization Technique

Consideration	Paravirtualization Technique	PCI Passthrough Technique
Interfaces	virtio (for KVM), VMXNET3 (for VMware)	SR-IOV
Use Cases	<ul style="list-style-type: none"> • Network simulation • Low-throughput applications 	<ul style="list-style-type: none"> • Static vMX deployments • High-throughput applications
Host Requirements	No requirements specific to this technique	Physical NIC must support PCI passthrough
VM Mobility (Junos OS Release 15.1F4 or earlier releases)	Moving vMX instance to a new server without reconfiguration.	Creating an identical vMX instance on a new server.

Paravirtualization

Supported in Junos OS Release 15.1F4, in a paravirtualized router, the VM and the host work together to efficiently move packets from the physical NIC to the application in the VM. You implement paravirtualization on the vMX router by configuring virtio, a technique that the KVM hypervisor supports that optimizes network and disk operations for the VM. Both the VFP VM and the host contain virtio drivers that interact to move packets. You implement paravirtualization on the VMware server by configuring VMXNET3 on the ESXi hypervisor. You must provide the following information in the configuration file for each vMX interface:

- Junos OS name
- Unique MAC address

If you want to move the VM from one server to another, you can do so without reconfiguration, provided the names and MAC addresses of each interface remain the same.

PCI Passthrough with SR-IOV

Supported in Junos OS Release 15.1F4, The vMX router supports PCI passthrough in combination with single root I/O virtualization (SR-IOV). In the PCI passthrough technique, you directly assign a NIC's memory space to a VM, enabling packets to bypass the hypervisor. Bypassing the hypervisor increases efficiency and results in high throughput of packets.

With SR-IOV, the hypervisor detects the physical NICs (known as a physical functions) and creates multiple virtual NICs (known as virtual functions) in the VFP VM. In the vMX implementation, the host dedicates a NIC to a single VM.

When you configure PCI passthrough with SR-IOV, you specify the following parameters for each vMX interface:

- Junos OS name
- Unique MAC address
- Name of the physical NIC

Because you create a direct connection between a virtual NIC and a physical NIC, you cannot move a VM from one host to another. If you need to move a VM to another host, you must install a new vMX instance on that host, and delete the vMX instance on the original host.

RELATED DOCUMENTATION

[vMX Overview | 2](#)

[Licenses for vMX](#)

2

CHAPTER

Prepare vMX Installation on VMWare

Minimum Hardware and Software Requirements | 11

vMX Package Contents | 15

Minimum Hardware and Software Requirements

IN THIS SECTION

- [Minimum Hardware Requirements for VMware | 11](#)
- [Software Requirements for VMware | 14](#)

This topic includes the following sections:

Minimum Hardware Requirements for VMware

[Table 2 on page 12](#) lists the hardware requirements.

Table 2: Minimum Hardware Requirements for VMware

Description	Value
<p>Number of cores</p> <p>NOTE: Performance mode is the default mode and the minimum value is based on one port.</p>	<p>For performance mode with low-bandwidth (virtio) or high-bandwidth (SR-IOV) applications: Minimum of 9</p> <ul style="list-style-type: none"> • 1 for VCP • 8 for VFP <p>The exact number of required vCPUs differs depending on the Junos OS features that are configured and other factors, such as average packet size. You can contact Juniper Networks Technical Assistance Center (JTAC) for validation of your configuration and make sure to test the full configuration under load before use in production. For typical configurations, we recommend the following formula to calculate the minimum vCPUs required by the VFP:</p> <ul style="list-style-type: none"> • Without QoS—$(4 * \textit{number-of-ports}) + 4$ • With QoS—$(5 * \textit{number-of-ports}) + 4$ <p>NOTE: All VFP vCPUs must be in the same physical non-uniform memory access (NUMA) node for optimal performance.</p> <p>In addition to vCPUs for the VFP, we recommend 2 x vCPUs for VCP and 2 x vCPUs for Host OS on any server running the vMX.</p>
	<p>For lite mode: Minimum of 4</p> <ul style="list-style-type: none"> • 1 for VCP • 3 for VFP <p>NOTE: If you want to use lite mode when you are running with more than 3 vCPUs for the VFP, you must explicitly configure lite mode.</p>

Table 2: Minimum Hardware Requirements for VMware (Continued)

Description	Value
<p>Memory</p> <p>NOTE: Performance mode is the default mode.</p>	<p>For performance mode:</p> <ul style="list-style-type: none"> • Minimum of 5 GB <ul style="list-style-type: none"> 1 GB for VCP 4 GB for VFP • Recommended of 16 GB <ul style="list-style-type: none"> 4 GB for VCP 12 GB for VFP <p>For lite mode: Minimum of 3 GB</p> <ul style="list-style-type: none"> • 1 GB for VCP • 2 GB for VFP
<p>Storage</p>	<p>Local or NAS</p> <p>Each vMX instance requires 44 GB of disk storage</p> <p>Minimum storage requirements:</p> <p>40 GB for VCP</p> <p>4 GB for VFP</p>

Table 2: Minimum Hardware Requirements for VMware (Continued)

Description	Value
vNICs	<ul style="list-style-type: none"> • SR-IOV <p>NOTE: SR-IOV is only supported with Intel Ivy Bridge CPU (or higher) and Intel x520 NICs using ixgbe driver or X710 NICs with 10G ports and using i40e driver. Any other NIC models are not supported.</p> <p>Support for unmodified ixgbe driver and i40e driver is available from Junos OS Release 18.4R1 onwards.</p> <ul style="list-style-type: none"> • VMXNET3 <p>NOTE: See "Virtual Network Interfaces for vMX" on page 7 for an overview on the types of virtual network interfaces supported by vMX.</p>

Software Requirements for VMware

[Table 3 on page 15](#) lists the software requirements.

Table 3: Software Requirements for VMware

Description	Value
Hypervisor	<ul style="list-style-type: none"> VMware ESXi 5.5 (Update 2) VMware ESXi 6.0 and 6.5 (Starting in Junos OS 18.1R1) VMware ESXi 6.7 and 7.0 (Starting in Junos OS 19.4R3) <p>NOTE: Due to a DPDK version change in Junos OS Release 18.1R1 and later, ESXi 6.5 is the minimum version required to run the vMX router if you are operating in high-bandwidth mode (performance mode). If you are operating the vMX router in low-bandwidth (lite) mode, you can use ESXi 6.0 or ESXi 5.5.</p> <p>NOTE: For performance mode, the minimum software requirement for ESXi 6.5 is Junos OS Release 18.1R1 and later.</p>
Management Client	vSphere 5.5 or vCenter Server

RELATED DOCUMENTATION[vMX Overview | 2](#)[Virtual Network Interfaces for vMX | 7](#)[Install vMX on VMWare | 18](#)[Manage vMX on VMWare | 28](#)

vMX Package Contents

[Table 4 on page 16](#) lists the images in the vMX package.

Table 4: vMX Package Contents

Filename	Description
<code>ova/vcp_*.ova</code>	Software image file for VCP.
<code>ova/vfpc_*.ova</code>	Software image file for VFP.

RELATED DOCUMENTATION

[Install vMX on VMWare | 18](#)

[Manage vMX on VMWare | 28](#)

[Minimum Hardware and Software Requirements | 11](#)

3

CHAPTER

Installing vMX

[Install vMX on VMWare](#) | 18

[Manage vMX on VMWare](#) | 28

[Enabling SR-IOV for VMware](#) | 33

Install vMX on VMWare

IN THIS SECTION

- [Installing vMX with OVA Files | 18](#)
- [Launching vMX on VMware | 28](#)

Read this topic to understand how to install vMX with the VMware vSphere Web Client using OVA files.

Installing vMX with OVA Files

IN THIS SECTION

- [Setting Up the Datastore | 19](#)
- [Setting Up the Network | 20](#)
- [Deploying the VCP VM | 24](#)
- [Deploying the VFP VM | 25](#)

vMX supports different physical PCI devices such as SR-IOV and VMXNET3. The procedure in this section is specific to VMXNET3 devices. See "[Enabling SR-IOV for VMware](#)" on page 33 for information to bring up vMX with SR-IOV NICs.

NOTE: If you are running vMX in performance mode, you must change the VM hardware version to 10.

To install vMX with the VMware vSphere Web Client using OVA files:

1. Download the vMX software package for VMware from the [vMX page](#) and uncompress the package in a location accessible to the server.

2. Launch the vSphere Web Client for your ESXi server and log in to the server with your credentials.

NOTE: You must set up these three interfaces to launch the VFP.

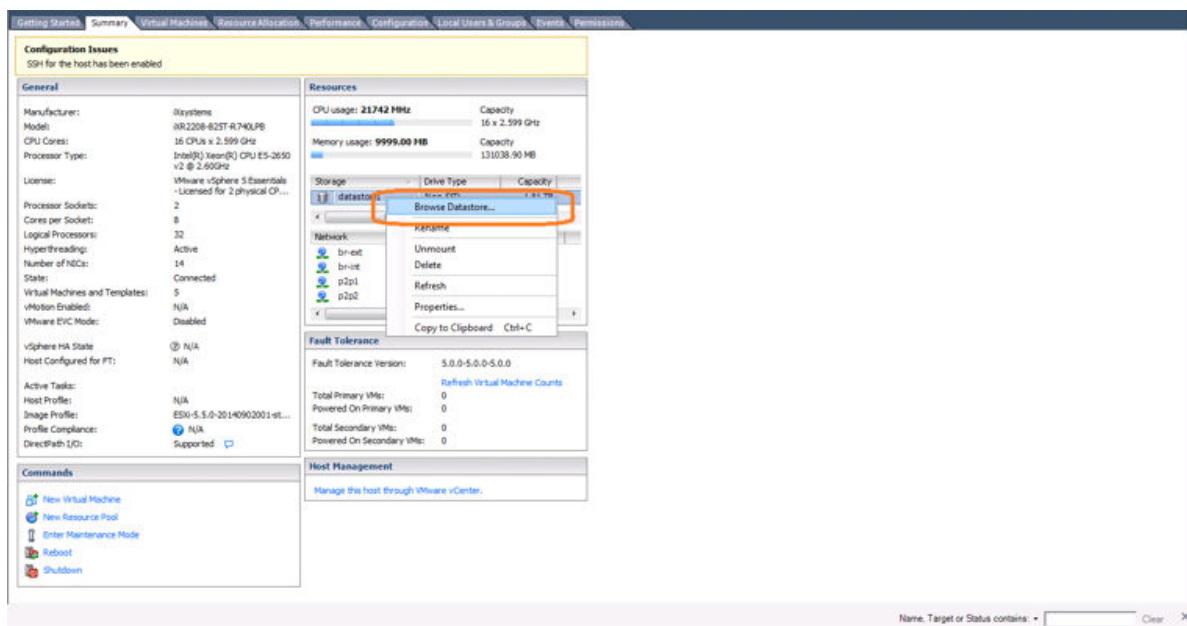
- Management access
- Bridge for internal communication between the VCP and VFP
- WAN interface (minimum of one)

To install vMX with vSphere for VMXNET3 adapters, perform these tasks:


Setting Up the Datastore

To upload vMX to the ESXi datastore:

1. Click the **Summary** tab, select the datastore under Storage, right-click, and select **Browse Datastore**.



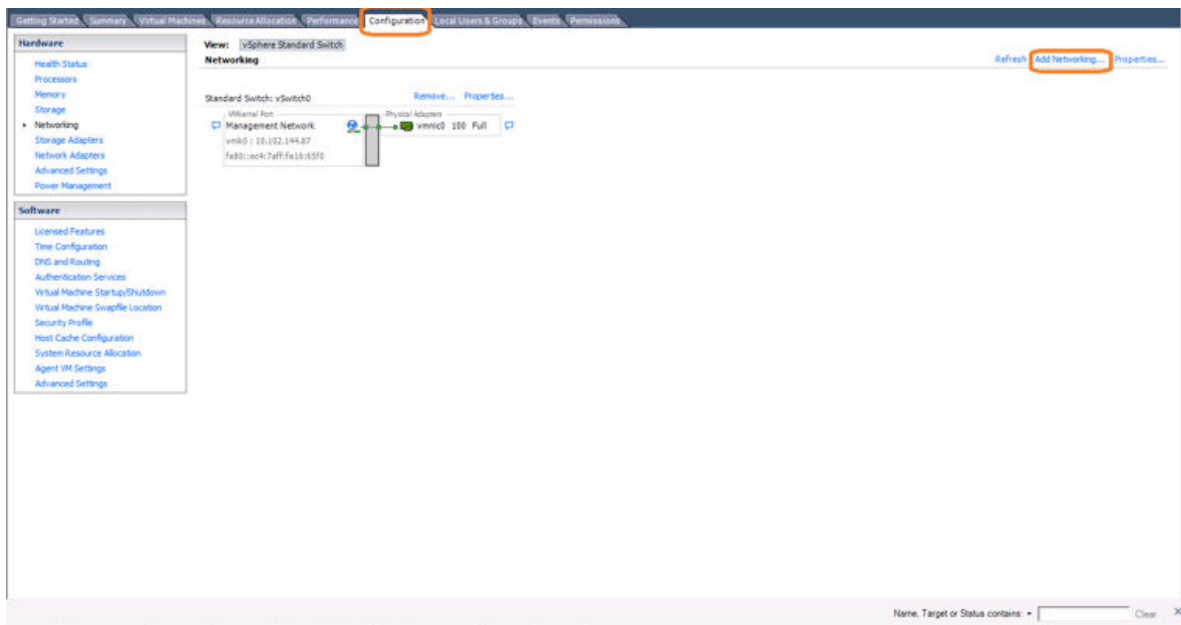
2. In the Datastore Browser, click the **Upload** button, select **Upload File**, and upload the files for the package contents listed in [vMX Package Contents on page 16](#).



Setting Up the Network

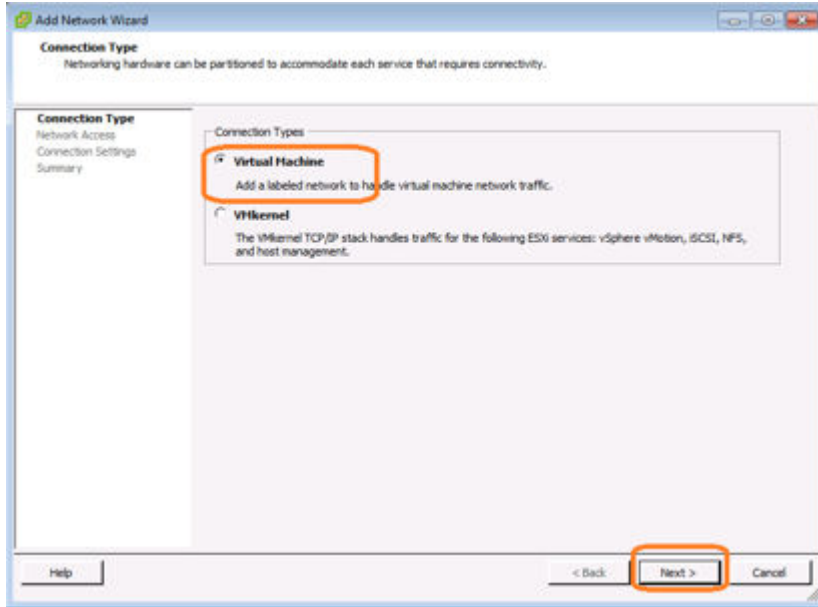
To set up the different networks for management (br-ext), internal connection of the VMs (br-int), and WAN ports for data:

1. In the left navigation pane, select the ESXi server and click the **Configuration** tab. Select **Networking** under Hardware.
2. In the top right corner, click **Add networking** to create the management vSwitch (br-ext).

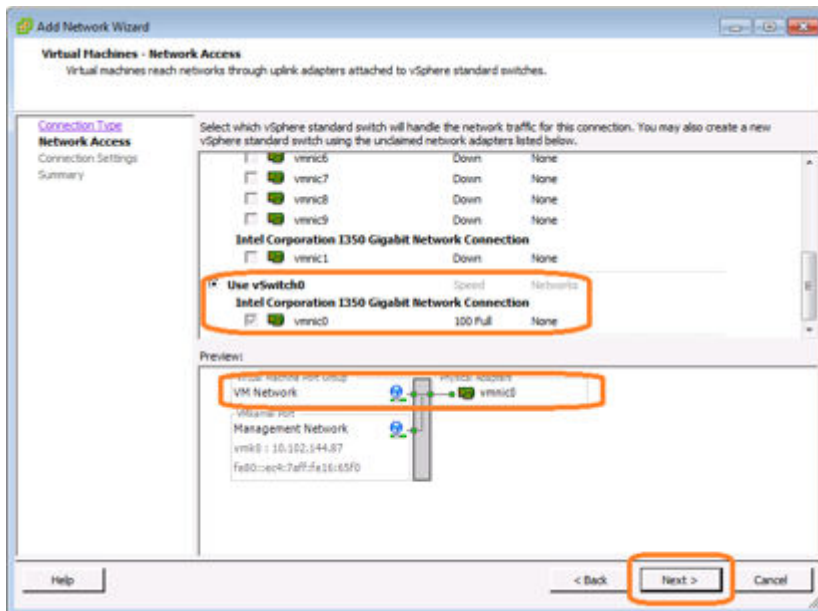


In the Add Network Wizard dialog box:

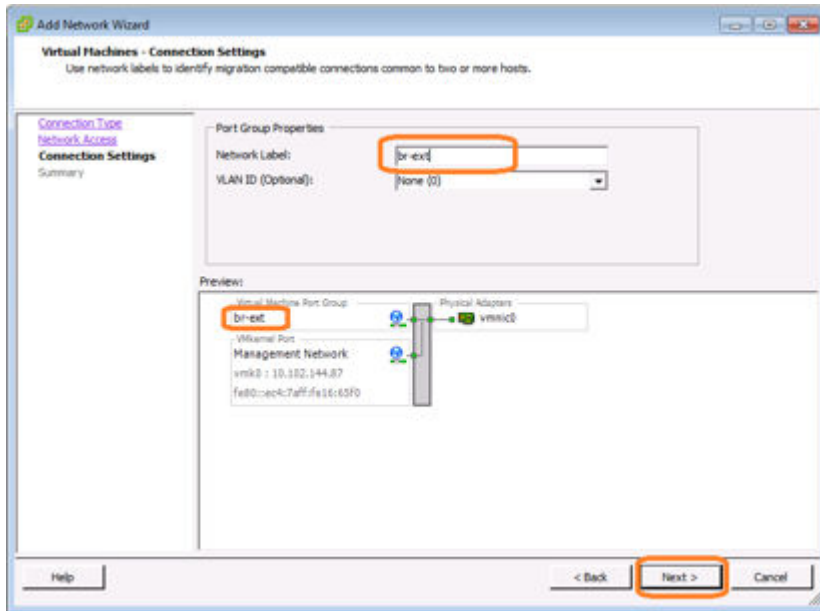
- a. For Connection Type, select the **Virtual Machine** option button and click **Next**.



- b. For Network Access, scroll down and select the **Use vSwitch0** option button, and click **Next**.

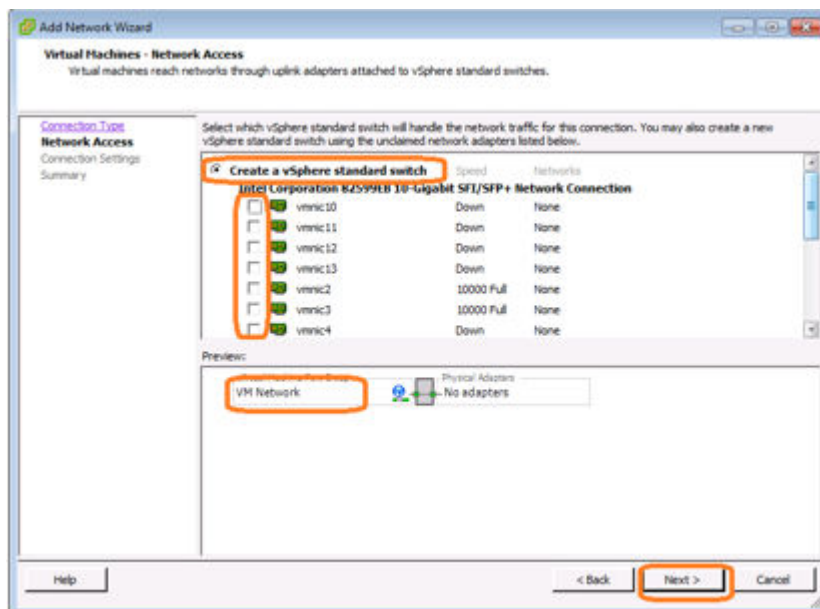


- c. For Connection Settings, name the network **br-ext** and click **Next**.



This network is connected to the management port (for example, vmnic0), which has a management IP address.

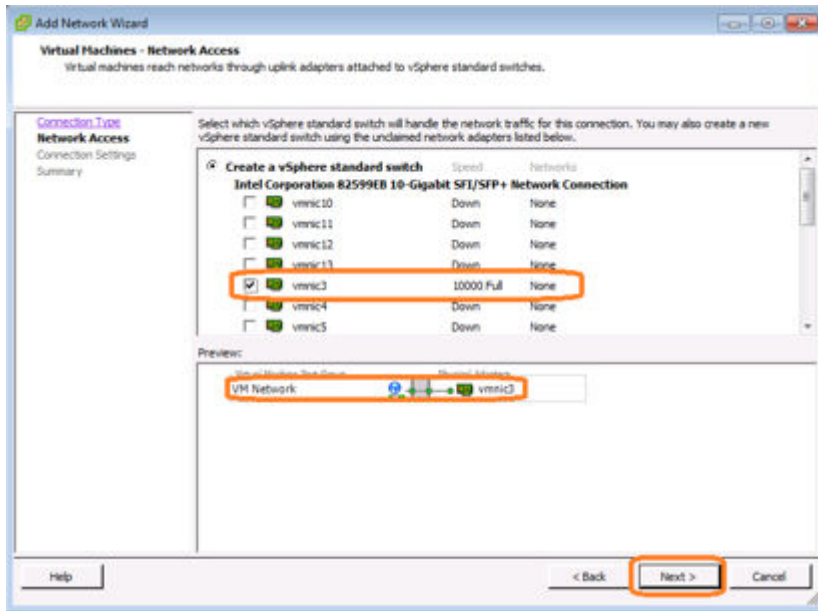
3. In the top right corner, click **Add networking** to add another vSwitch (br-int). In the Add Network Wizard dialog box:
 - a. For Connection Type, select the **Create vSphere standard switch** option button, clear all the physical NIC check boxes (vmnicx), and click **Next**.



- b. For Connection Settings, name the network **br-int** and click **Next**.

Note that this network is not connected to any physical NICs.

4. In the top right corner, click **Add networking** to add another vSwitch for the WAN port (p2p1).
 - a. For Connection Type, select the **Create vSphere standard switch** option button, select the physical NIC check box to which you want to direct WAN traffic (for example, vmnic2), and click **Next**.
The link is up as indicated by 10000Full.

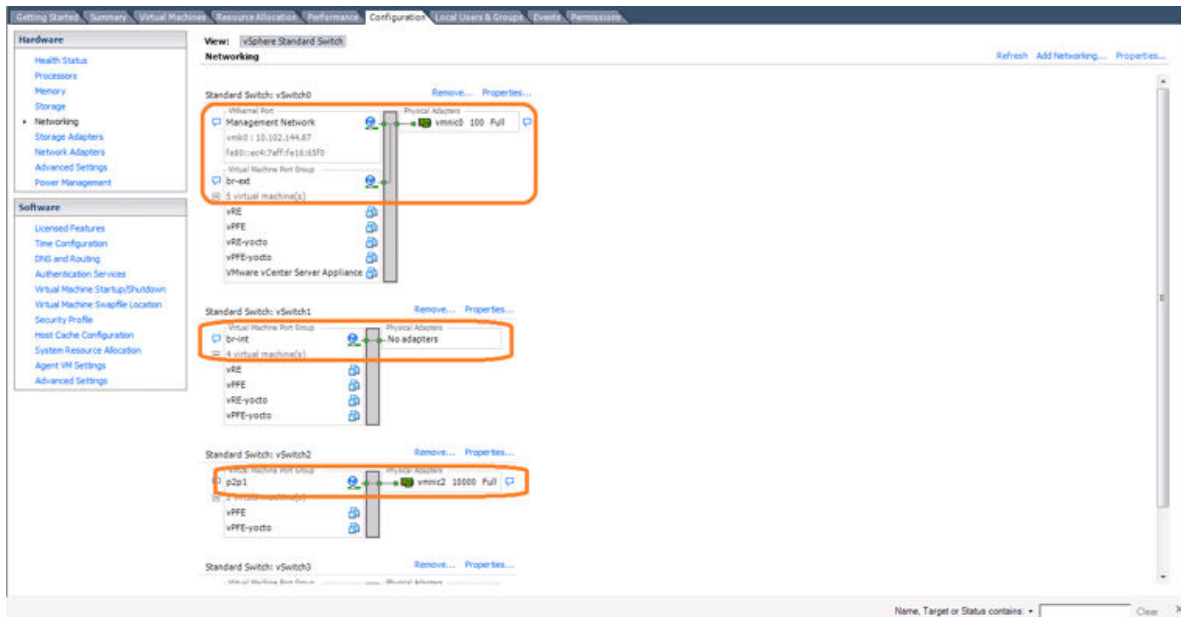


- b. For Connection Settings, name the network **p2p1**.
5. In the top right corner, click **Add networking** to add another vSwitch for the WAN port (p2p2).
 - a. For Connection Type, select the **Create vSphere standard switch** option button, select the physical NIC check box to which you want to direct WAN traffic (for example, vmnic3), and click **Next**.
The link is up as indicated by 10000Full.
 - b. For Connection Settings, name the network **p2p2**.

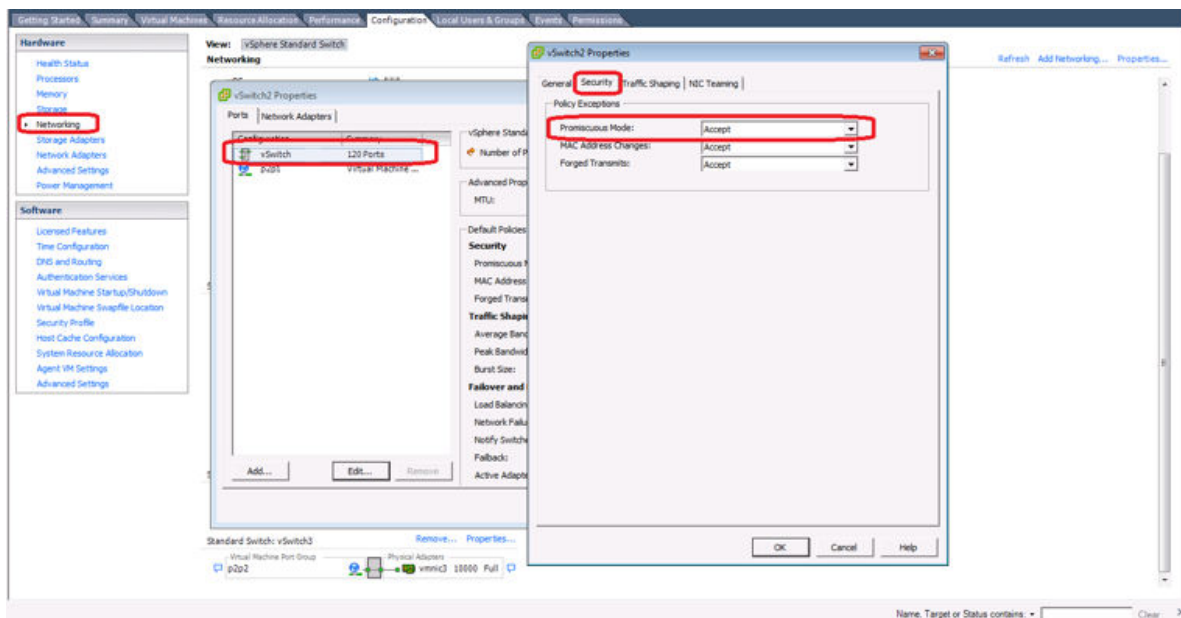
You can see the four vSwitches you created (br-ext, br-int, p2p1, p2p2) in the Configuration tab for Networking. You must enable promiscuous mode in all vSwitches so that packets with any MAC addresses can reach the vMX. For example, OSPF needs multicast MAC address communication so you must enable promiscuous mode to support it.

To enable promiscuous mode in a vSwitch:

1. In the Configuration tab for Networking, click **Properties** next to the Standard Switch for each vSwitch.



2. In the vSwitch Properties dialog box, select vSwitch in the Ports tab and click **Edit** at the bottom.
3. In the vSwitch Properties dialog box, select **Accept** from the Promiscuous Mode list in the Security tab and click **OK**.



Deploying the VCP VM

To deploy the VCP VM using .ova files:

1. Select the host and select **Deploy OVF Template** from the **File** menu.
2. In the Source pane, click **Browse**, select the .ova file for the VCP, and click **Next**.

3. In the OVF Template Details pane, click **Next**. This pane displays a summary of the OVA file contents.
4. In the Name and Location pane, specify the name of the VM and click **Next**.
5. In the Storage pane, select the appropriate datastore for the destination storage of the VM and click **Next**.
6. In the Disk Format pane, select the **Thick Provision Lazy Zeroed** option button and click **Next**.
7. In the Network Mapping pane, map the destination network on the host to the source network. For VCP, you must map the source networks for the external management network for connecting to management interfaces (br-ext) and the internal connection between VCP and VFP (br-int).
Select the destination network (for example, br-ext) for the br-ext source network, select the destination network (for example, br-int) for the br-int source network, and click **Next**.
8. In the Ready to Complete pane, verify your configuration and click **Finish**.

The **.ova** file is deployed as the VCP VM.

Deploying the VFP VM

To deploy the VFP VM using **.ova** files:

1. Select the host and select **Deploy OVF Template** from the **File** menu.
2. In the Source pane, click **Browse**, select the **.ova** file for the VFP, and click **Next**.
3. In the OVF Template Details pane, click **Next**. This pane displays a summary of the OVA file contents.
4. In the Name and Location pane, specify the name of the VM and click **Next**.
5. In the Storage pane, select the appropriate datastore for the destination storage of the VM and click **Next**.
6. In the Disk Format pane, select the **Thick Provision Lazy Zeroed** option button and click **Next**.
7. In the Network Mapping pane, map the destination network on the host to the source network. For VFP, you must map the source networks for the external management network for connecting to management interfaces (br-ext) and the internal connection between VCP and VFP (br-int).
Select the destination network (for example, br-ext) for the br-ext source network, select the destination network (for example, br-int) for the br-int source network, and click **Next**.
8. In the Ready to Complete pane, verify your configuration and click **Finish**.

The **.ova** file is deployed as the VFP VM.

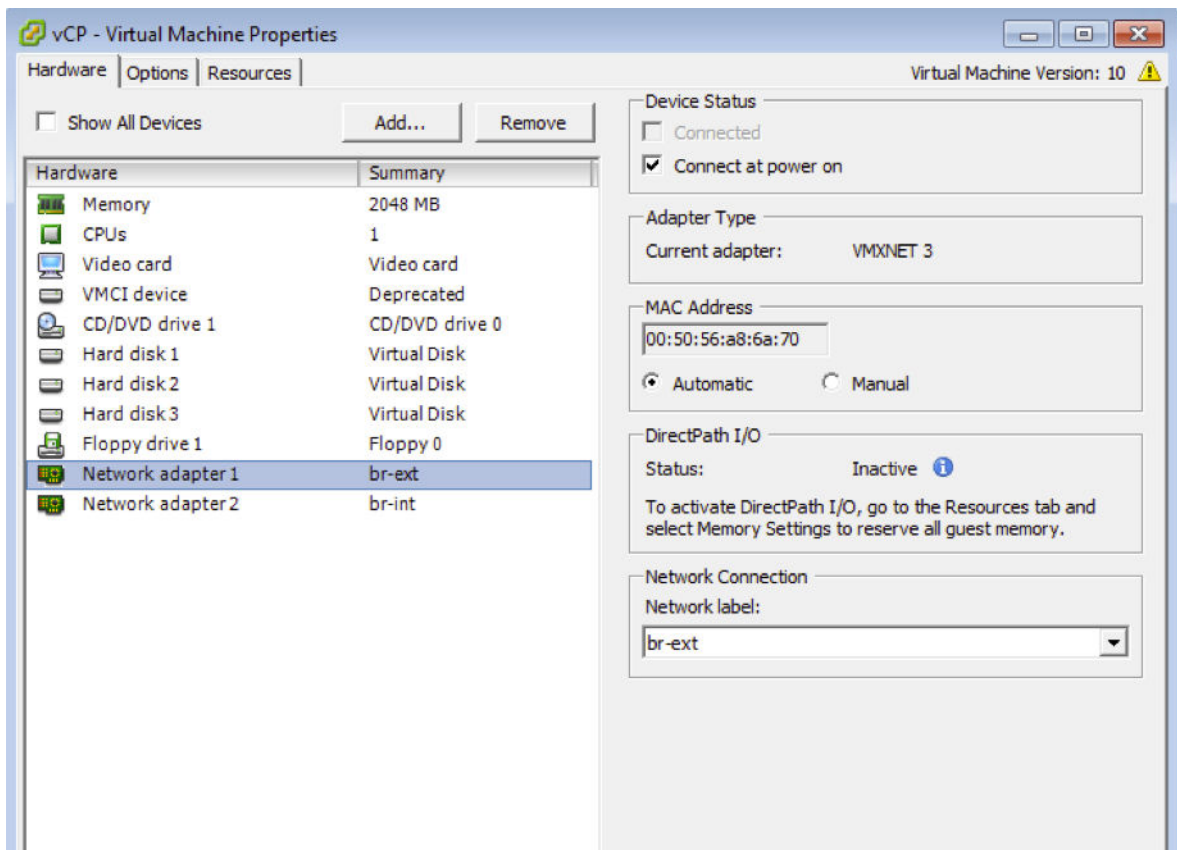
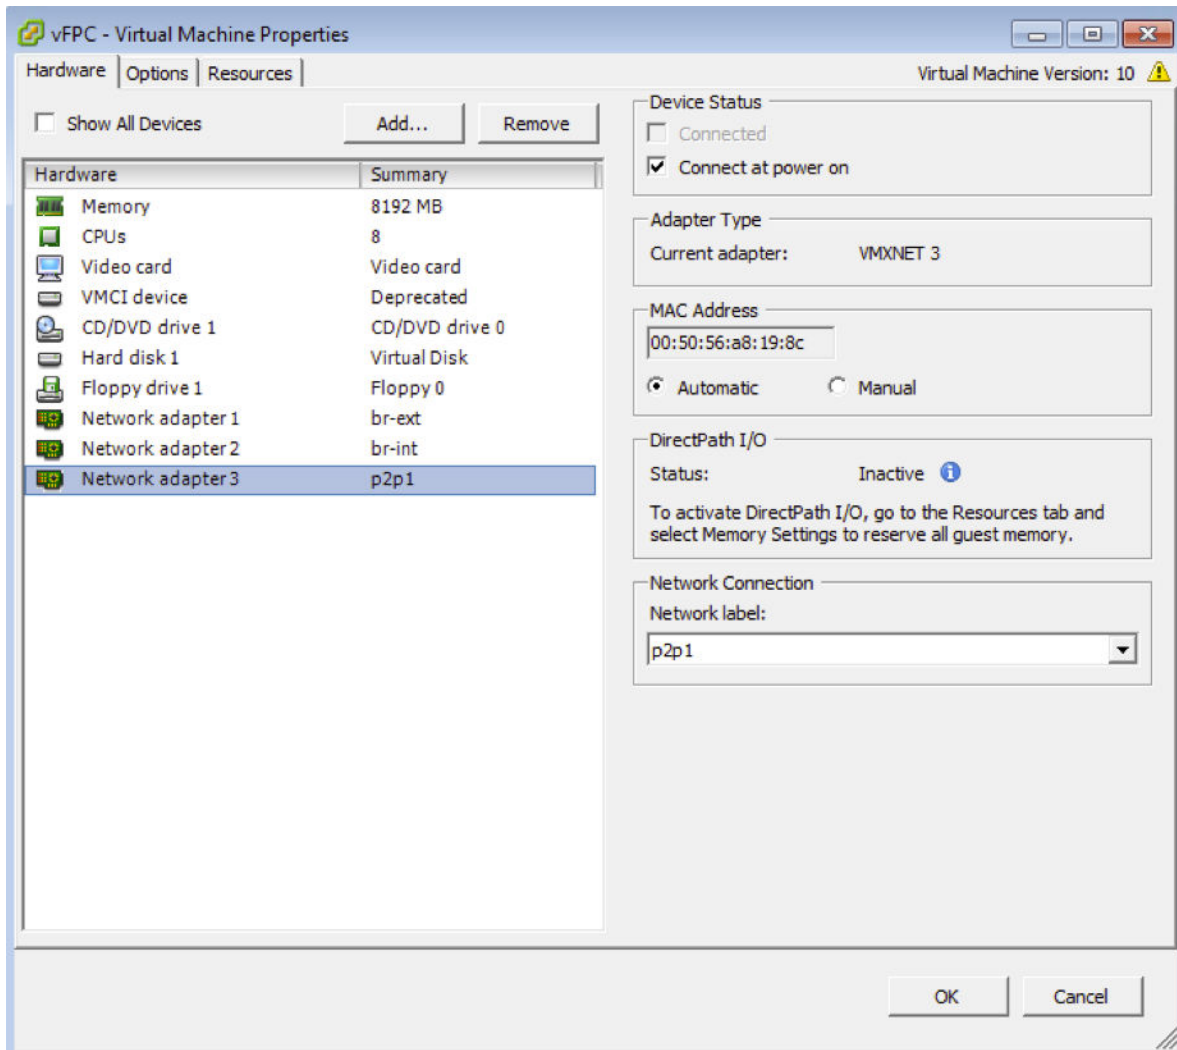
After you have deployed the VFP VM, you can modify the amount of memory, the number of vCPUs, and the number of WAN ports.

NOTE: Before you launch the VFP VM, make sure you have configured the proper number of vCPUs and memory for your VM based on the requirements described in "[Minimum Hardware and Software Requirements](#)" on page 11.

To modify these settings for the VFP VM:

1. In the left navigation pane, select the VM and right-click **Edit Settings** to display the Virtual Machine Properties window.
2. To change the amount of memory, select Memory in the Hardware tab of the Virtual Machine Properties window. Change the memory size and click **OK**.
3. To change the number of vCPUs, select CPUs in the Hardware tab of the Virtual Machine Properties window. Change the number of virtual sockets and click **OK**.
4. To add WAN ports, click **Add** in the Hardware tab of the Virtual Machine Properties window. The Add Hardware wizard is displayed.
 - a. For Device Type, select **Ethernet Adapter** and click **Next**.
 - b. For Network connection, select the adapter type (for example, VMXNET3), select the network connection (for example, p2p1 or p2p2), and click **Next**.
 - c. Verify your configuration and click **Finish**.

Click **OK** in the Virtual Machine Properties window.



Launching vMX on VMware

After you power on both VMs, vMX is started as a virtual network function (VNF).

To launch vMX:

1. Select the VFP VM in the left navigation pane and right-click **Power > Power On**.
2. Select the VCP VM in the left navigation pane and right-click **Power > Power On**.

NOTE: You must shut down the vMX instance before you reboot host server using the request system halt command.

RELATED DOCUMENTATION

[Minimum Hardware and Software Requirements | 11](#)

[vMX Package Contents | 15](#)

[Manage vMX on VMWare | 28](#)

Manage vMX on VMWare

IN THIS SECTION

- [Allowing VLAN Traffic | 29](#)
- [Setting Up a Serial Port Connection to the VM | 29](#)
- [Viewing CPU Information | 32](#)
- [Verifying Whether VMs Are Running | 32](#)

Read this topic to understand how to perform tasks such as allowing VLAN traffic and setting up serial port connection post vMX installation.

Allowing VLAN Traffic

Note the following regarding VLAN usage on vMX ESXi:

- vMX on VMware ESXi supports 64 VLANs per IFD.
- For VLANs to work on the VMware ESXi, you must set VLAN offload on the NICs that are being used.

To allow VLAN traffic to reach vMX, configure the interface to allow all VLAN IDs.

To configure the interface to pass all VLAN traffic:

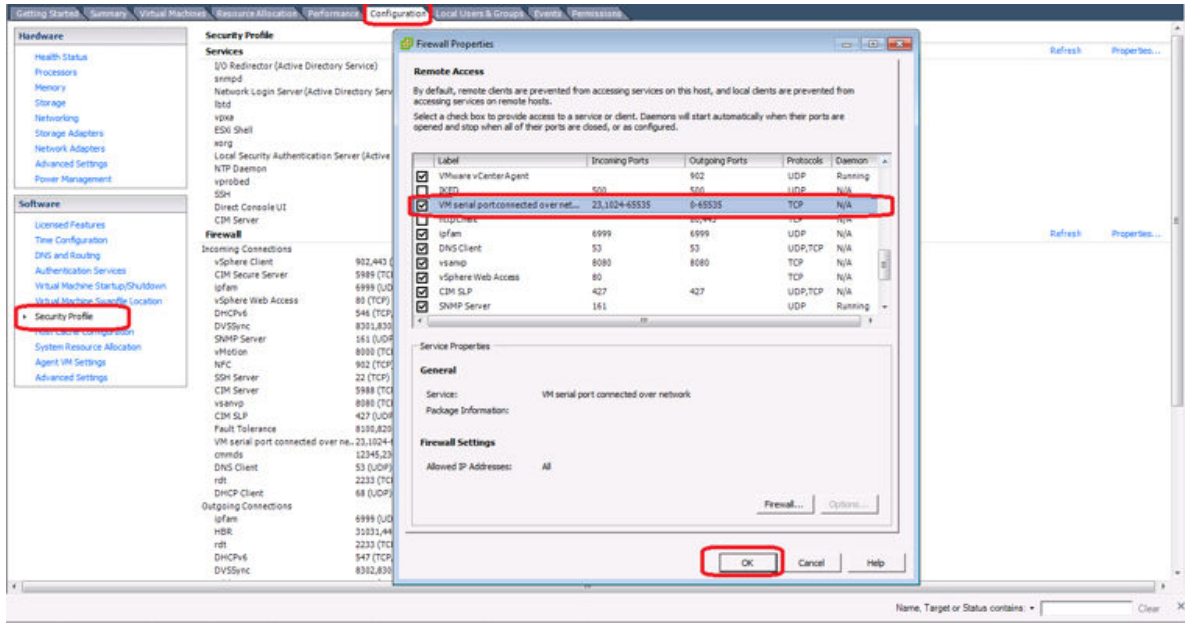
1. In the left navigation pane, select the ESXi server and click the **Configuration** tab.
2. Select **Networking** and click **Properties** next to the Standard Switch that contains your WAN interface.
3. In the vSwitch Properties dialog box, select the network interface in the Port tab and click **Edit** at the bottom.
4. In the Port Group Properties section, select **All** from the VLAN ID list in the General tab and click **OK**.

Setting Up a Serial Port Connection to the VM

You can set up the ESXi firewall to allow incoming serial port connections so that you can access the VM using telnet.

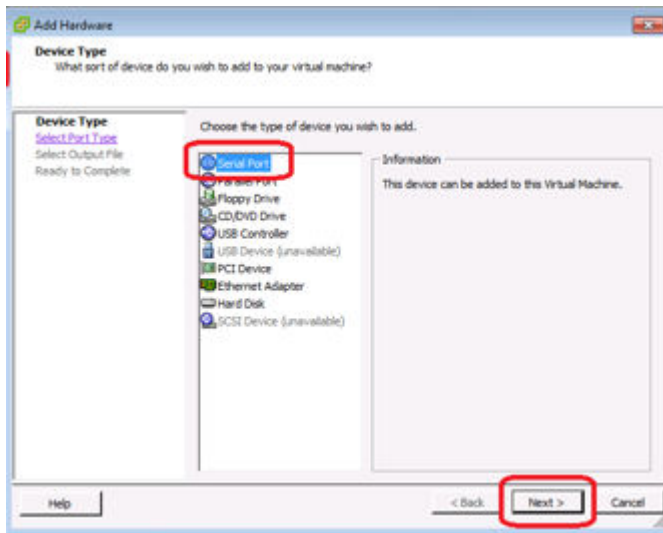
To configure the serial port connection for a VM:

1. In the left navigation pane, select the ESXi server and click the **Configuration** tab.
2. Select Security Profile and click **Properties** next to Firewall.
3. In the Firewall Properties dialog box, select the **VM serial port connected over network** box and click **OK**.

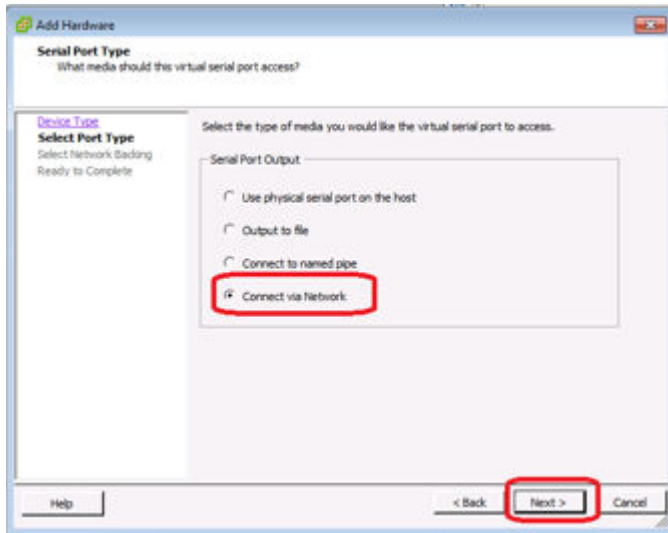


4. In the left navigation pane, select the VM and right-click **Edit Settings** to display the Virtual Machine Properties window. In the Hardware tab of the Virtual Machine Properties window, click **Add**. The Add Hardware wizard is displayed.

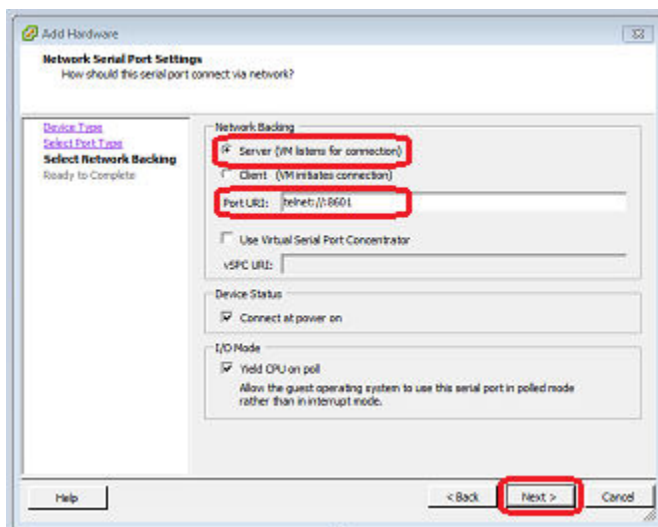
a. For Device Type, select **Serial Port** and click **Next**.



b. For Select Port Type, select the **Connect via Network** button and click **Next**.



- c. For Select Network Backing, select the **Server** button and specify the port number in the Port URI text box in the format **telnet://:port-number** (for example, **telnet://:8601**). Click **Next**.



- d. Click **Finish**.

After you have configured the serial port connection, you can access the VM using the **telnet esxi-server-ip-address port-number** command.

Logging In to VCP

You can access the serial console using the **telnet esxi-server-ip-address port-number** command, and log in with the username **root** and no password.

To disconnect from the console, log out of the session and press Ctrl +]. At the **telnet>** prompt, type **close** and press Enter.

Logging In to VFP

You can access the serial console using the **telnet *esxi-server-ip-address port-number*** command, and log in with the username **root** and password **root**.

To disconnect from the console, log out of the session and press Ctrl +]. At the **telnet>** prompt, type **close** and press Enter.

Viewing CPU Information

To display CPU information, select the ESXi server in the left navigation pane and click the **Summary** tab in the Inventory view.

Verifying Whether VMs Are Running

To verify that the VMs are running after vMX is installed, select the VM in the left navigation pane and click the **Summary** tab of the VM. The Summary tab displays the current status and whether the VM is powered on or off.

RELATED DOCUMENTATION

[Install vMX on VMWare | 18](#)

[Minimum Hardware and Software Requirements | 11](#)

[vMX Package Contents | 15](#)

Enabling SR-IOV for VMware

IN THIS SECTION

- [Enable SR-IOV on a Physical Adapter | 33](#)
- [Add vSwitch and Port Groups | 35](#)
- [Assigning the SR-IOV NIC to the VFP VM | 36](#)
- [Enable VLAN Tagging on SR-IOV Interfaces | 39](#)
- [SR-IOV Interface for Layer 2 Services | 41](#)

The physical network interfaces support single root I/O virtualization (*SR-IOV*) capability and can be connected to the VMs using PCI passthrough. Before you enable SR-IOV for VMware, note the following:

- SR-IOV is only supported with Intel Ivy Bridge CPU (or higher) and Intel X520 NICs using ixgbe driver or X710 NICs with 10G ports and using i40e driver. Any other NIC models are not supported.
Support for ixgbe driver and i40e driver is available from Junos OS Release 18.4R1 onwards.
- Starting in Junos OS Release 18.4R1, in VMware deployments operating in SR-IOV mode with an ESXi server, support is available for VLAN-tagged traffic for vMX interfaces.

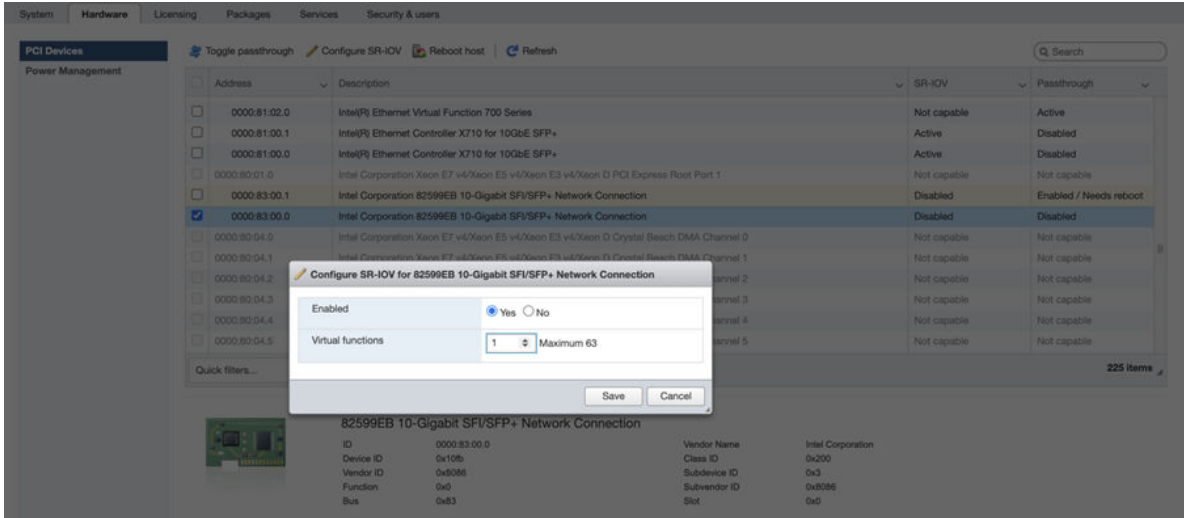
To enable vMX with vSphere for SR-IOV adapters, perform these tasks:

Enable SR-IOV on a Physical Adapter

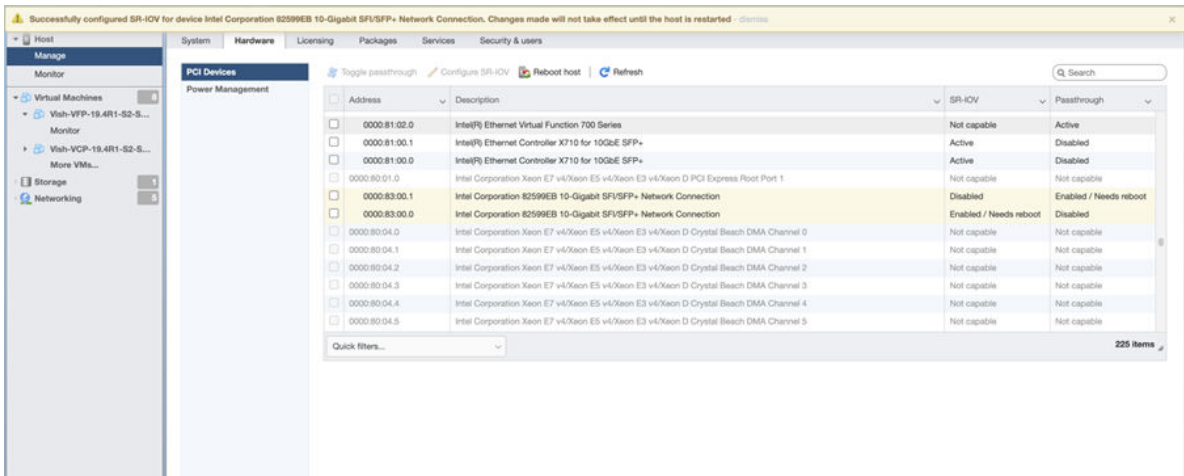
Use the following steps to enable SR-IOV on a physical adapter. The following procedures are applicable for ESXi version 6.5.

To enable SR-IOV on a physical adapter in vSphere client:

1. Navigate to the host in the left navigation pane and click the **Manage** tab. Go to **Hardware** tab. Under **PCI Devices** option, select the physical adapter and click **Configure SR-IOV** option.

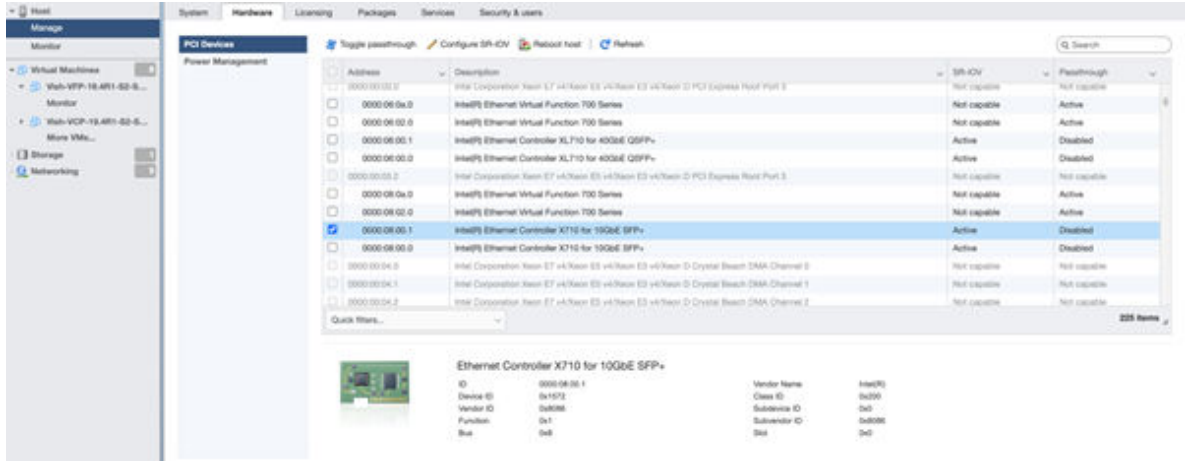


2. In the populated window, select **Yes** to enable SR-IOV. In the virtual functions text box, specify the number of virtual functions to configure for the adapter.



3. Click **Save**.
4. Restart the host.

5. Verify the status of SR-IOV post reboot. The status is displayed as Active.

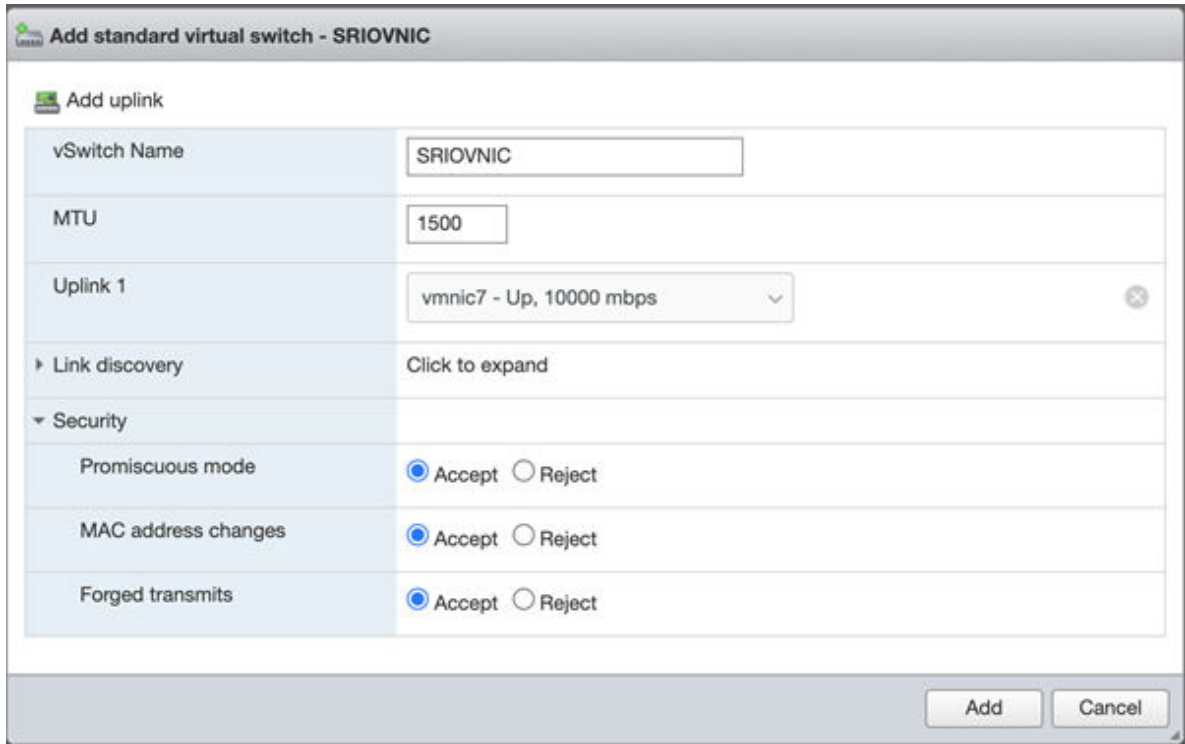


Add vSwitch and Port Groups

Use the following steps to enable SRX-IOV on a physical adapter if you are using ESXi version 6.5:

To add vSwitch and Port Groups:

1. To create vSwitch: Navigate to **Networking** tab and select **Virtual Switches** and click **Add standard virtual switch**.



2. To create port groups: Navigate to **Networking** tab and select **Port Groups** and click **Add port group**.

Add port group - SRIOVNIC-PortGroup	
Name	SRIOVNIC-PortGroup
VLAN ID	0
Virtual switch	SRIOVNIC
▼ Security	
Promiscuous mode	<input type="radio"/> Accept <input type="radio"/> Reject <input checked="" type="radio"/> Inherit from vSwitch
MAC address changes	<input type="radio"/> Accept <input type="radio"/> Reject <input checked="" type="radio"/> Inherit from vSwitch
Forged transmits	<input type="radio"/> Accept <input type="radio"/> Reject <input checked="" type="radio"/> Inherit from vSwitch
<input type="button" value="Add"/> <input type="button" value="Cancel"/>	

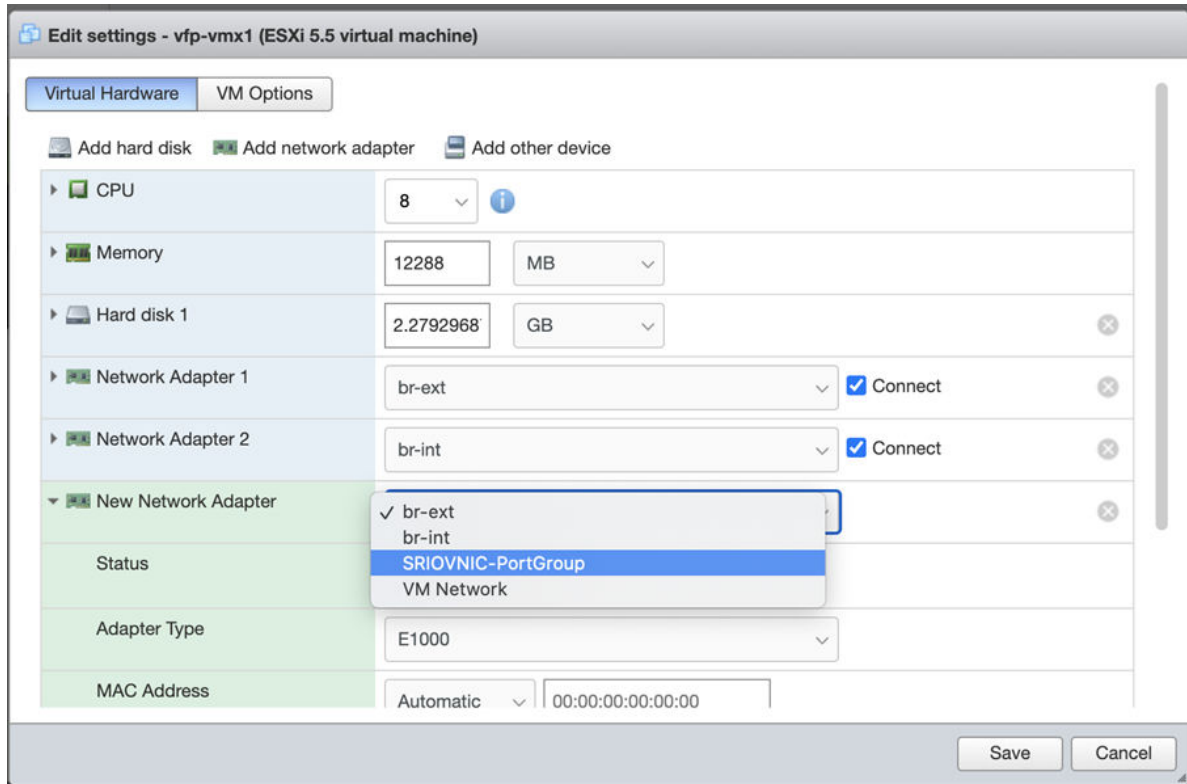
Assigning the SR-IOV NIC to the VFP VM

Use the following steps to enable SRX-IOV on a physical adapter if you are using ESXi version 6.5:

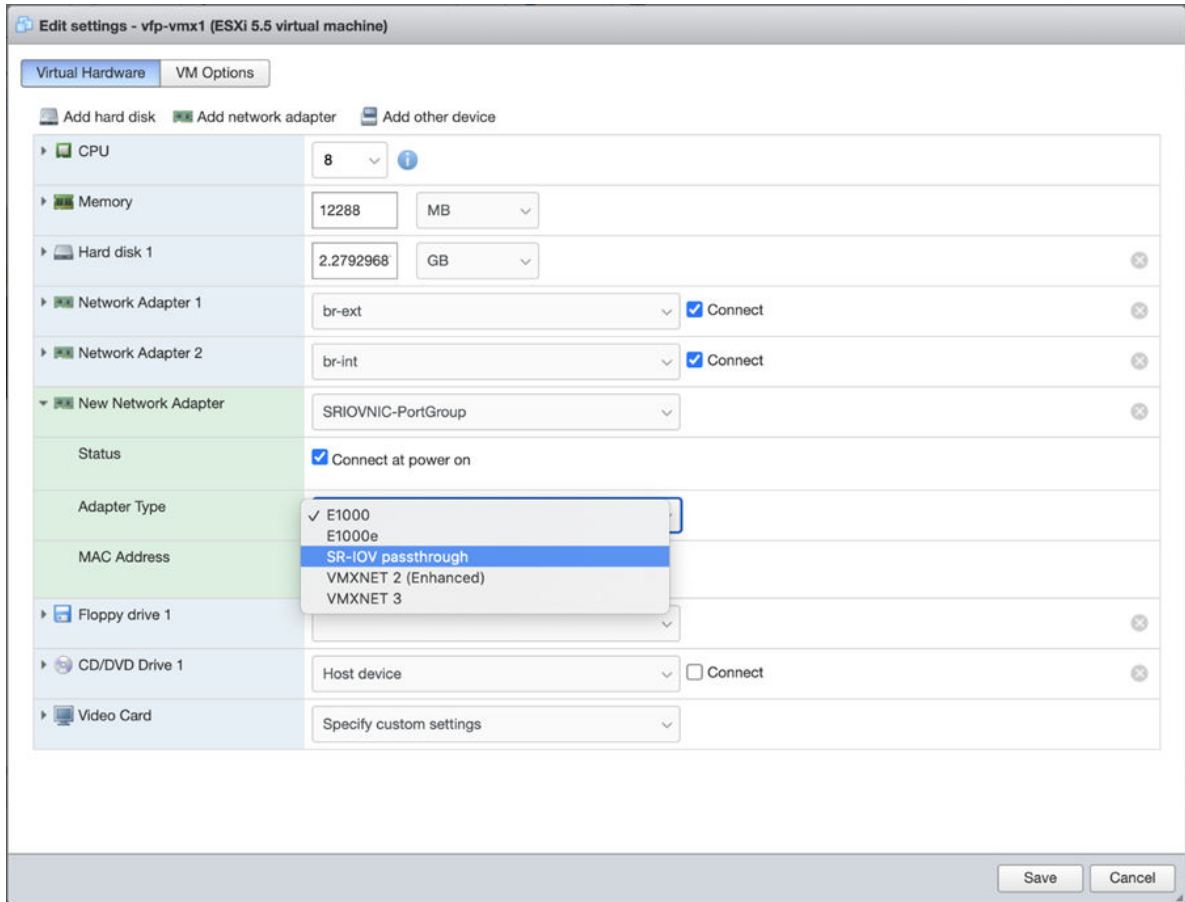
To assign the SR-IOV to the VFP VM using the vSphere Web Client:

1. Power-off the VFP when you add SR-IOV to the VFP VM.
2. Navigate to the VFP VM in the left navigation pane and click the **Manage** tab.
3. Select **VM Hardware**, and click **Edit** near the top right corner.
4. In the Edit setting, go to **Virtual Hardware** tab and click **Add network adapter** option.

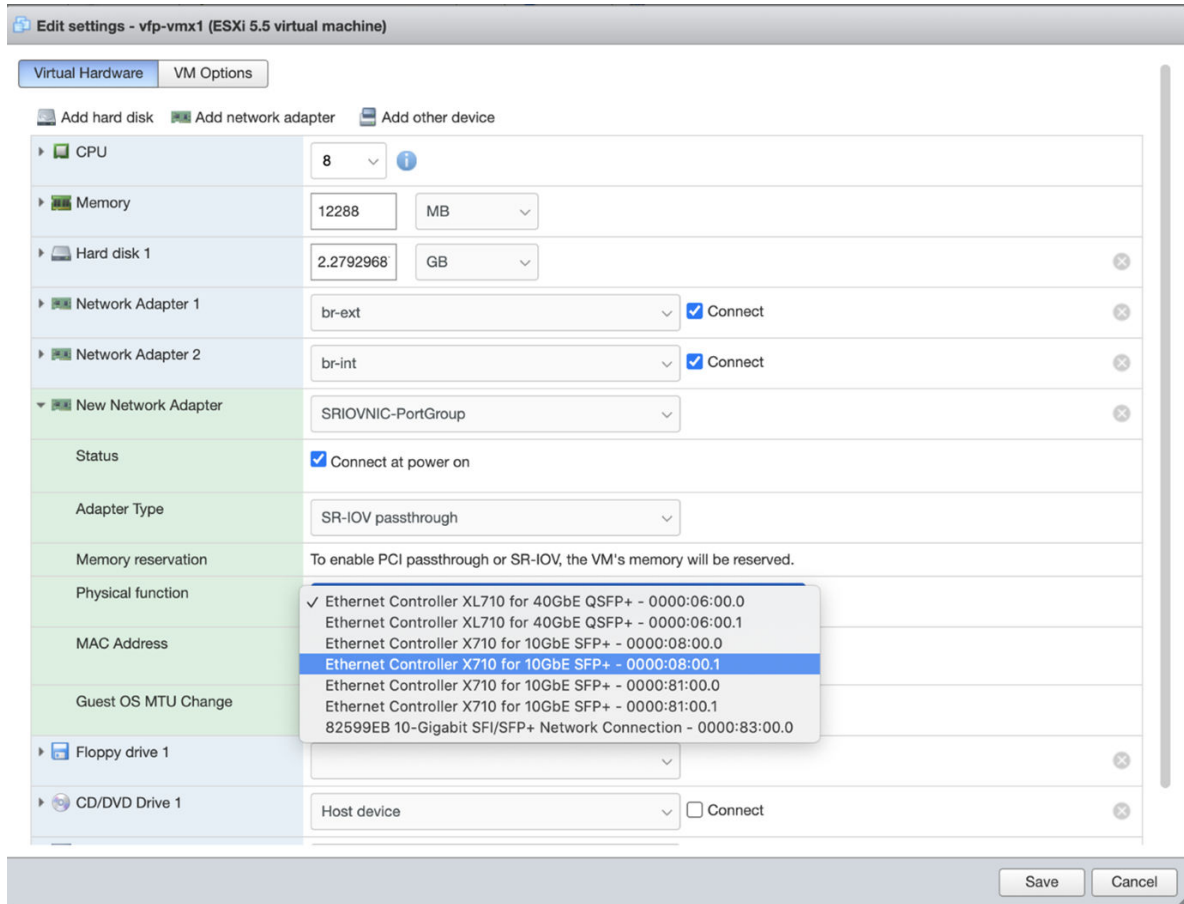
5. Under New Network Adapter, select the respective port group for mapping.



6. Change **Adapter Type** as SR-IOV passthrough.



7. Map **Physical function** to respective SR-IOV NIC.



8. Click **Save**.

Now, the SR-IOV interface is added to the VFP VM.

Enable VLAN Tagging on SR-IOV Interfaces

You can configure VLAN tagging on SR-IOV interfaces on vMX instance deployed in VMware ESXi host with the following steps::

We've used the hardware and software combination for the following procedure:

- VMware ESXi version 6.5 Update 3
- i40en driver version 1.10.6 for Intel 700 Series network adapters
- Firmware version 7.20 for Intel 700 Series network adapters

1. Enter the CLI configuration mode after logging in to the vMX.

2. Include the **vlan-tagging** statement to enable VLAN tagging.

```
[edit interfaces xe-0/0/0]
user@host# set vlan-tagging
```

3. Include the **vlan-offload** statement to offload the VLAN filtering .

```
[edit interfaces xe-0/0/0]
user@host# set vlan-offload
```

Example:

```
user@host> show configuration interfaces xe-0/0/0
vlan-tagging;
vlan-offload;
encapsulation flexible-ethernet-services;
unit 0 {
    encapsulation vlan-vpls;
    vlan-id 10;
    family vpls;
}
unit 1 {
    vlan-id 30;
    family inet {
        address 30.30.30.1/30;
    }
}
unit 2 {
    encapsulation vlan-bridge;
    vlan-id 40;
}
```

4. Log in to the ESXi Web console . Navigate to **Networking** and select **Port groups** tab. Edit the settings for the required port.

Edit port group - SRIOV-PG-vmnic12	
Name	SRIOV-PG-vmnic12
VLAN ID	4095
Virtual switch	SRIOV-vmnic12
▸ Security	Click to expand
▸ NIC teaming	Click to expand
▸ Traffic shaping	Click to expand

Save Cancel

Enable the VLAN trunking for the vSwitch that includes the SR-IOV interface. You must set the VLAN value to 4095.

Note the following:

- The VLAN ID 4095 allows any VLAN traffic.
- The VLAN ID 0 does not allow any VLAN traffic.
- Ensure that you have configured the VLAN ID matching to the VLAN configuration on the mapped interface. Example: If you have configured a VLAN ID 100 on the associated logical interface, then use VLAN ID 100 for SR-IOV. This setting enables receive and forwarding incoming frames that contain a matching VLAN ID that is VLAN 100.

5. Click **Save** to save your settings.

SR-IOV Interface for Layer 2 Services

To use SR-IOV interface for Layer 2 Services such as VPLS, Layer 2 VPN, Layer 2 circuit or bridging on a vMX instance deployed in VMware ESXi host, you must enable the trust mode and disable the spoof check in Intel NIC settings. Use the following steps to enable the trust mode and disable the spoof check:

1. Install Intel esxcli plug-in. For instructions, see [VMware KB article](#).
2. Verify the trust mode and the spoof check status on the SR-IOV NIC by entering the following command in ESXi shell mode:

```
esxcli intnet sriovnic vf get -v <vf number> -n <vmnic name>
```

Example:

```
root@host:~] esxcli intnet sriovnic vf get -v 0 -n vmnic12
```

VF ID	Trusted	Spoof Check
0	false	true

3. Enable the trust mode and disable the spoof check on Intel NIC.

```
esxcli intnet sriovnic vf set -s false -t true -v <vf number> -n <vmnic name>
```

Example:

```
[root@host:~] esxcli intnet sriovnic vf set -s false -t true -v 0 -n vmnic12
```

Trusted mode is set to true and spoof check is set to false

```
[root@host:~] esxcli intnet sriovnic vf get -v 0 -n vmnic12
```

VF ID	Trusted	Spoof Check
0	true	false

4. Restart the vMX FPC configured with the SR-IOV interface after changing to the trust mode and spoof check values.

```
user@host> request chassis fpc slot <number> restart
```

Once the FPC is online, Layer-2 services start working on SR-IOV interface.

RELATED DOCUMENTATION

[Virtual Network Interfaces for vMX | 7](#)

[Minimum Hardware and Software Requirements | 11](#)

[Install vMX on VMWare | 18](#)

4

CHAPTER

Configuring vMX Chassis-Level Features

[Configuring the Number of Active Ports on vMX | 45](#)

[Naming the Interfaces | 45](#)

[Configuring the Media MTU | 46](#)

[Enabling Performance Mode or Lite Mode | 47](#)

[Tuning Performance Mode | 49](#)

Configuring the Number of Active Ports on vMX

You can specify the number of active ports for vMX. The default number of ports is 10, but you can specify any value in the range of 1 through 23. You can change this number if you want to limit the number of Ethernet interfaces in the VCP VM to match the number of NICs added to the VFP VM.

NOTE: If you are running virtio interfaces in lite mode, you can use up to 96 ports. Other configurations running in performance mode support up to 23 ports.

To specify the number of active ports, configure the number of ports at the `[edit chassis fpc 0 pic 0]` hierarchy level.

```
[edit]
user@vmx# set chassis fpc 0 pic 0 number-of-ports
```

RELATED DOCUMENTATION

[Naming the Interfaces | 45](#)

[Configuring the Media MTU | 46](#)

[Enabling Performance Mode or Lite Mode | 47](#)

[Tuning Performance Mode | 49](#)

Naming the Interfaces

vMX supports the following interface types:

- Gigabit Ethernet (ge)
- 10-Gigabit Ethernet (xe)
- 100-Gigabit Ethernet (et)

By default, the interfaces come up as ge interfaces with 1 Gbps bandwidth in the Junos OS configuration. The default port speed values for the interface types are 1 Gbps (ge), 10 Gbps (xe), and 100 Gbps (et). If you do not enable schedulers, the speed is only for display purposes and is not

enforced. If you enable schedulers, the transmit rate of the port is limited to the speed unless it is overridden by the shaping rate in the CoS configuration.

To specify the interface types, configure the interface type at the `[edit chassis fpc 0 pic 0]` hierarchy level.

```
[edit]
user@vmx# set chassis fpc 0 pic 0 interface-type (ge | xe | et)
```

RELATED DOCUMENTATION

[Configuring the Number of Active Ports on vMX | 45](#)

[Configuring the Media MTU | 46](#)

[Enabling Performance Mode or Lite Mode | 47](#)

[Tuning Performance Mode | 49](#)

Configuring the Media MTU

For vMX, you can configure the media MTU in the range 256 through 9500.

NOTE: For VMware, the maximum value is 9000. For AWS, the maximum value is 1514.

You configure the MTU by including the `mtu` statement at the `[edit interface interface-name]` hierarchy level.

```
[edit]
user@vmx# set interface ge-0/0/0 mtu bytes
```

RELATED DOCUMENTATION

[Configuring the Number of Active Ports on vMX | 45](#)

[Naming the Interfaces | 45](#)

Enabling Performance Mode or Lite Mode

VMX can be configured to run in two modes depending on the use case.

- Lite mode—Needs fewer resources in terms of CPU and memory to run at lower bandwidth.
- Performance mode—Needs higher resources in terms of CPU and memory to run at higher bandwidth.

NOTE: Starting in Junos OS Release 15.1F6 and later releases performance mode is enabled implicitly by default.

When you enable performance mode, make sure you have configured the proper number of vCPUs (four or more VPCUs) and memory for your VMs based on your use case.

You can explicitly enable lite-mode. If you are using paravirtualized network interfaces such as virtio (for KVM) or VMXNET3 (for VMware) for lab simulation use cases, you can disable performance mode by including the **lite-mode** statement at the **[edit chassis fpc 0]** hierarchy level.

```
[edit]
user@vmx# set chassis fpc 0 lite-mode
```

You can explicitly enable performance mode by including the **performance-mode** statement at the **[edit chassis fpc 0]** hierarchy level.

```
[edit]
user@vmx# set chassis fpc 0 performance-mode
```

NOTE: We recommend that you enable hyperthreading in BIOS. We recommend that you verify the process with the vendor because different systems have different methods to enable hyperthreading.

Starting with Junos OS Release 17.3R1, the **show chassis hardware** command displays the mode in which vMX is running in the part number field for the FPC. RIOT-PERF indicates performance mode and RIOT-LITE indicates lite mode. For example, this output indicates that vMX is running in lite mode.

```
user@vmx> show chassis hardware
```

```
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis                               VM54599D128A  VMX
Midplane
Routing Engine 0                        RE-VMX
CB 0                                     VMX SCB
CB 1                                     VMX SCB
FPC 0                                     Virtual FPC
  CPU          Rev. 1.0 RIOT-LITE  BUILTIN
  MIC 0
  PIC 0                BUILTIN  BUILTIN  Virtual
```

[Table 5 on page 48](#) highlights some of the challenging features which are supported in the Fast Path and some which are not supported. Features which are not supported in the Fast Path still work but they get less than 100K PPS per worker vCPU.

Table 5: Features Support in Fast Path

Features	Support in Fast Path
Pseudowire Headend Termination (PWHT) (Layer 2 VPN)	Not Supported
L2 circuit	Not Supported
Ethernet VPN (EVPN)	Not Supported

Table 5: Features Support in Fast Path (Continued)

Features	Support in Fast Path
Virtual Extensible LAN protocol (VXLAN)	Not Supported
MPLS-over-UDP (MPLSoUDP)	Not Supported
Inline J-flow	Supported
Pseudowire Headend Termination (PWHT) (Layer 3 VPN and IP)	Supported
GRE	Supported
logical tunnel interfaces (lt)	Supported

Release History Table

Release	Description
15.1F6	Starting in Junos OS Release 15.1F6 and later releases performance mode is enabled implicitly by default.

RELATED DOCUMENTATION

[Tuning Performance Mode | 49](#)

[lite-mode | 0](#)

[performance-mode | 0](#)

Tuning Performance Mode

To tune performance mode for the traffic, you can specify the number of Workers dedicated to processing multicast and control traffic. You can specify any value in the range of 0 through 15. The default of 0 specifies that all available Workers are used to process all traffic.

The number of dedicated Workers specified in relation to the number of available Workers results in the following behavior:

- If the number of dedicated Workers is greater than or equal to the number of available Workers, then all available Workers are used to process all traffic.
- If the number of dedicated Workers is less than the number of available Workers, then the first set of available Workers (equal to the specified number of dedicated Workers) is used to process multicast and control traffic while the remaining available Workers are used to process flow cache traffic.

To specify the number of dedicated Workers for processing multicast and control traffic, configure the number of Workers at the `[edit chassis fpc 0 performance-mode]` hierarchy level.

```
[edit]  
user@vmx# set chassis fpc 0 performance-mode number-of-ucode-workers number-workers
```

NOTE: Changing the number of Workers reboots the FPC.

RELATED DOCUMENTATION

[Enabling Performance Mode or Lite Mode | 47](#)

[performance-mode | 0](#)

5

CHAPTER

Class of Service for vMX

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)

[Configuring Hierarchical CoS on vMX | 56](#)

[Example: Configuring Hierarchical CoS on vMX | 59](#)

[Configuring Four-Level Hierarchical Scheduling on vMX | 65](#)

[Packet Loss Priority and Drop Profiles on vMX | 66](#)

[Managing Congestion Using Drop Profiles and Packet Loss Priorities on vMX | 68](#)

CoS on vMX Overview

vMX supports two-level hierarchical scheduling (per-unit scheduler or hierarchical scheduler) with VLAN queuing. Each VLAN (logical interface) uses three traffic classes and eight queues.

Starting with Junos OS Release 17.3R1, vMX supports four-level hierarchical scheduling for up to 16 level 2 CoS scheduler nodes. The level 2 node maps to the interface set or VLAN (logical interface).

vMX supports shaping at the traffic class level, not at the queue level. A traffic class is a bundle of queues with fixed priority. The next level in the hierarchy is the VLAN (logical interface), which is a bundle of traffic classes.

vMX has the following fixed priorities and queues for these traffic classes:

- Traffic Class 1: High (strict priority)

Queue 0

Queue 6

- Traffic Class 2: Medium (strict priority)

Queue 1

Queue 7

- Traffic Class 3: Low

Queue 2

Queue 3

Queue 4

Queue 5

NOTE: Both Traffic Class 1 and Traffic Class 2 follow strict priority, so all excess traffic is discarded as tail drops. However, Traffic Class 3 does not follow strict priority, so the shaping rate is set to the shaping rate of the VLAN.

All queues in the same traffic class have equal priority, so the scheduler pulls packets from each queue in the traffic class based on weighted round robin (WRR) for the VLAN.

All configured forwarding classes must be mapped to one of the queues.

The following features are not supported::

- Weighted random early detection (WRED)
- Queue buffer size configuration

NOTE: No commit errors are displayed for unsupported features.

Starting in Junos OS Release 18.4R1, the quality of service (QoS) configuration is enhanced such that, when a port is oversubscribed and congested, a subscriber with higher priority gets more weight than a subscriber with a lower priority. For example, when a subscriber on a port has 100 MB service and another subscriber has 10 MB service then the subscriber with 100 MB service gets more priority than the subscriber with 10 MB service. You must ensure that the priority is followed at level 1 and level 2 nodes, regardless of the weight. The WRR provides the ability handle the oversubscription so that the scheduled traffic reflects a ratio of the shaping rate configured for the individual VLANs.

Use the following commands to configure a maximum number of 16384 subscribers per port on a level 2 node and a maximum number of 32768 subscribers per port on a level 3 node:

```
set interfaces <interface-name> hierarchical-scheduler maximum-hierarchy 3 max-
l2-nodes 16384
set interfaces <interface-name> hierarchical-scheduler maximum-hierarchy 3 max-
l3-nodes 32768
```

NOTE: The default number of subscribers that are configured per level 2 node is 4000.

Use the following command to disable the WRR feature:

```
subport_oversubscription_disable=1 in the /etc/riot/runtime.conf of the vFP
```

The following list describes the limitations for WRR:

- The delay-buffer rate must be configured for WRR to work appropriately.
- A discrepancy in the delay-buffer rate values, among the VLANs belonging to the same level 2 scheduler node can cause the WRR to work incorrectly.
- The WRR works incorrectly when the ratio of shaping rate is greater than 100 among all the subscribers.
- The number of level 2 scheduler nodes and the number of subscribers per level 2 scheduler node must be equal to 32,000.

- Any modification to the level 2 scheduler node configuration would require a FPC reset.

RELATED DOCUMENTATION

[CoS Features and Limitations on vMX | 54](#)

[Packet Loss Priority and Drop Profiles on vMX | 66](#)

CoS Features and Limitations on vMX

IN THIS SECTION

- [Weighted Round-Robin of Subscriber Traffic on a Port Limitations | 55](#)

vMX has the following limitations for CoS support:

- Schedulers support only the **transmit-rate** and **excess-rate** statements. Only weights are supported at the queue level, so transmission rate and excess rate are used for calculating queue weights.
 - If **transmit-rate percent** is configured at the queue level, then configure guaranteed rate at the VLAN level.

NOTE: Guaranteed rate is not supported, but it is used to calculate queue weights.

- If you only configure transmit rate, queue weights are calculated based on the transmission rate.
- If you only configure excess rate, queue weights are calculated based on the excess rate.
- If you configure both transmit rate and excess rate, queue weights are calculated based on the excess rate.
- If you configure the excess rate for one queue, the excess rate is expected for all the queues to compute the weights. If the excess rate is not configured, the default weight of 1 is used.

NOTE: To get the expected behavior, you must configure the excess rate for all queues.

- Traffic control profiles support only the **shaping-rate** and **scheduler-map** statements.
If a traffic control profile has a default scheduler map, you must configure the guaranteed rate.
- For high- and medium-priority traffic classes, the transmission rate is the shaping rate.
- For low-priority queues, the shaping rate for the VLAN is used for the queue. As a result, the low-priority queues can burst up to the configured shaping rate for the VLAN. The transmission rate is used as the WRR weight when there is more than one queue configured for a given priority.

Some considerations for the high- and medium-priority traffic classes:

- All excess traffic from the traffic classes for high- and medium-priority queues are discarded as tail drops.
- For high- and medium-priority traffic classes, the transmission rate is the shaping rate.
If the transmission rate is not configured and the shaping rate is configured, then the queue weight is calculated based upon the configured shaping rate.
If you configure the transmission rate for both queues of the same traffic class, the shaping rate of the traffic class is the sum of the individual transmission rates of the queues for that traffic class.
- If a queue is not configured, its transmission rate is set to zero.
If no queues are configured, the shaping rate of the VLAN is applied to the traffic class as the transmission rate.
- If any of the queues in the traffic class is configured, the shaping rate of the VLAN is set to the guaranteed rate of the configured queue. If a queue is not configured, the guaranteed rate is set to zero by default.
- If the sum of the rates of the individual queues in a traffic class exceeds the shaping rate of the VLAN, the shaping rate of the VLAN is used as the shaping rate of the traffic class.

Weighted Round-Robin of Subscriber Traffic on a Port Limitations

The following list describes the limitations for WRR:

- A discrepancy in the delay-buffer rate values among the VLANs belonging to the same level 2 scheduler node can cause the WRR to work incorrectly.
- WRR does not work correctly if the ratio of the shaping rate is greater than 100 among all the subscribers.

- The number of level 2 scheduler nodes and the number of subscribers per level 2 scheduler node must be equal to 32,000 for it to work correctly.
- Any modification to the level 2 scheduler node configuration requires an FPC reset.

RELATED DOCUMENTATION

[Configuring Hierarchical CoS on vMX | 56](#)

[CoS on vMX Overview | 52](#)

Configuring Hierarchical CoS on vMX

IN THIS SECTION

- [Enabling Flexible Queuing | 56](#)
- [Mapping Forwarding Classes to Queues on vMX | 57](#)
- [Configuring Traffic Control Profiles for vMX | 57](#)
- [Configuring Schedulers on vMX | 57](#)

To configure hierarchical CoS, perform these tasks:

Enabling Flexible Queuing

Hierarchical CoS is disabled by default. To enable hierarchical CoS, include the **flexible-queuing-mode** statement at the **[edit chassis fpc 0]** hierarchy level and restart the FPC.

```
[edit]
user@vmx# set chassis fpc 0 flexible-queuing-mode
```

Mapping Forwarding Classes to Queues on vMX

You must map all configured forwarding classes to one of the queues.

```
[edit]
user@vmx# set class-of-service forwarding-classes class class-name queue-num queue-number
```

Configuring Traffic Control Profiles for vMX

Traffic control profiles support only the **shaping-rate** and **scheduler-map** statements for vMX.

To specify the shaping rate, include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level.

```
[edit]
user@vmx# set class-of-service traffic-control-profiles profile-name shaping-rate rate
```

To specify the scheduler map, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level.

```
[edit]
user@vmx# set class-of-service traffic-control-profiles profile-name scheduler-map map-name
```

Configuring Schedulers on vMX

The scheduler map contains the mapping of forwarding classes to their schedulers. The scheduler defines the properties for the queue.

Schedulers support only the **transmit-rate** and **excess-rate proportion** statements for vMX.

To specify the transmission rate, include the **transmit-rate** statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level.

```
[edit]
user@vmx# set class-of-service schedulers scheduler-name transmit-rate rate
```

BEST PRACTICE: Guaranteed rate is not supported, so there is no reserved bandwidth for the VLAN. To get the expected behavior, we recommend that you configure the transmit rate to be the guaranteed rate.

To specify the proportion of the excess bandwidth to share, include the **excess-rate proportion** statement at the [edit class-of-service schedulers *scheduler-name*] hierarchy level. The value is in the range of 0 through 1000.

```
[edit]
user@vmx# set class-of-service schedulers scheduler-name excess-rate proportion value
```

If you configure the excess rate for one queue, the excess rate is expected for all the queues to compute the weights. If the excess rate is not configured, the default weight of 1 is used.

NOTE: To get the expected behavior, you must configure the excess rate for all queues. For example, if you configure excess rate for the low-priority queues, configure the same excess rate for the high- and medium-priority queues.

RELATED DOCUMENTATION

[Example: Configuring Hierarchical CoS on vMX | 59](#)

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)

Example: Configuring Hierarchical CoS on vMX

IN THIS SECTION

- [Requirements | 59](#)
- [Overview | 59](#)
- [Configuration | 59](#)

This example describes how to configure hierarchical CoS on vMX with eight queues.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 16.2
- vMX Release 16.2

Overview

This example configures two-level hierarchical schedulers with specified transmission rates.

Configuration

IN THIS SECTION

- [Configuring the Chassis | 60](#)
- [Applying Shaping and Scheduling to VLANs | 60](#)

Configuring the Chassis

CLI Quick Configuration

```
[edit]
set chassis fpc 0 flexible-queuing-mode
```

Step-by-Step Procedure

To enable hierarchical CoS on the chassis:

1. Enable flexible queuing mode on the chassis.

```
[edit]
user@vmx# set chassis fpc 0 flexible-queuing-mode
```

Once you commit the configuration, the FPC is restarted.

Applying Shaping and Scheduling to VLANs

CLI Quick Configuration

```
[edit]
set class-of-service forwarding-classes class voice1 queue-num 0
set class-of-service forwarding-classes class video1 queue-num 1
set class-of-service forwarding-classes class data1 queue-num 2
set class-of-service forwarding-classes class data2 queue-num 3
set class-of-service forwarding-classes class data3 queue-num 4
set class-of-service forwarding-classes class data4 queue-num 5
set class-of-service forwarding-classes class voice2 queue-num 6
set class-of-service forwarding-classes class video2 queue-num 7
set interfaces ge-0/0/0 hierarchical-scheduler maximum-hierarchy-levels 2
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 100 vlan-id 100
set interfaces ge-0/0/0 unit 100 family inet address 10.2.2.1/24
set interfaces ge-0/0/1 hierarchical-scheduler maximum-hierarchy-levels 2
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 100 vlan-id 100
set interfaces ge-0/0/1 unit 100 family inet address 10.1.1.1/24
```

```
set class-of-service classifiers inet-precedence vlan_tos forwarding-class
voicel loss-priority low code-points 000
set class-of-service classifiers inet-precedence vlan_tos forwarding-class
videol loss-priority low code-points 001
set class-of-service classifiers inet-precedence vlan_tos forwarding-class data1
loss-priority low code-points 010
set class-of-service classifiers inet-precedence vlan_tos forwarding-class data2
loss-priority low code-points 011
set class-of-service classifiers inet-precedence vlan_tos forwarding-class data3
loss-priority low code-points 100
set class-of-service classifiers inet-precedence vlan_tos forwarding-class data4
loss-priority low code-points 101
set class-of-service classifiers inet-precedence vlan_tos forwarding-class
voice2 loss-priority low code-points 110
set class-of-service classifiers inet-precedence vlan_tos forwarding-class
video2 loss-priority low code-points 111
set class-of-service traffic-control-profiles ge_0_0_1_vlan_100_tcp shaping-rate
50m
set class-of-service traffic-control-profiles ge_0_0_1_vlan_100_tcp scheduler-
map vlan_smap
set class-of-service interfaces ge-0/0/1 unit 100 output-traffic-control-profile
ge_0_0_1_vlan_100_tcp
set class-of-service interfaces ge-0/0/0 unit 100 classifiers inet-precedence
vlan_tos
set class-of-service scheduler-maps vlan_smap forwarding-class voicel scheduler
sched_voicel
set class-of-service scheduler-maps vlan_smap forwarding-class videol scheduler
sched_video1
set class-of-service scheduler-maps vlan_smap forwarding-class data1 scheduler
sched_data1
set class-of-service scheduler-maps vlan_smap forwarding-class data2 scheduler
sched_data2
set class-of-service scheduler-maps vlan_smap forwarding-class data3 scheduler
sched_data3
set class-of-service scheduler-maps vlan_smap forwarding-class data4 scheduler
sched_data4
set class-of-service scheduler-maps vlan_smap forwarding-class voice2 scheduler
sched_voice2
set class-of-service scheduler-maps vlan_smap forwarding-class video2 scheduler
sched_video2
set class-of-service schedulers sched_voicel transmit-rate 15m
set class-of-service schedulers sched_video1 transmit-rate 15m
set class-of-service schedulers sched_data1 transmit-rate 5m
```

```

set class-of-service schedulers sched_data2 transmit-rate 5m
set class-of-service schedulers sched_data3 transmit-rate 5m
set class-of-service schedulers sched_data4 transmit-rate 5m
set class-of-service schedulers sched_voice2 transmit-rate 10m
set class-of-service schedulers sched_video2 transmit-rate 10m

```

Step-by-Step Procedure

To apply shaping and scheduling:

1. Map the forwarding classes to their respective queues.

```

[edit]
user@vmx# set class-of-service forwarding-classes class voice1 queue-num 0
user@vmx# set class-of-service forwarding-classes class video1 queue-num 1
user@vmx# set class-of-service forwarding-classes class data1 queue-num 2
user@vmx# set class-of-service forwarding-classes class data2 queue-num 3
user@vmx# set class-of-service forwarding-classes class data3 queue-num 4
user@vmx# set class-of-service forwarding-classes class data4 queue-num 5
user@vmx# set class-of-service forwarding-classes class voice2 queue-num 6
user@vmx# set class-of-service forwarding-classes class video2 queue-num 7

```

2. Configure the interfaces to enable two-level hierarchical scheduling and apply scheduling to the VLANs.

```

[edit]
user@vmx# set interfaces ge-0/0/0 hierarchical-scheduler maximum-hierarchy-levels 2
user@vmx# set interfaces ge-0/0/0 vlan-tagging
user@vmx# set interfaces ge-0/0/0 unit 100 vlan-id 100
user@vmx# set interfaces ge-0/0/0 unit 100 family inet address 10.2.2.1/24
user@vmx# set interfaces ge-0/0/1 hierarchical-scheduler maximum-hierarchy-levels 2
user@vmx# set interfaces ge-0/0/1 vlan-tagging
user@vmx# set interfaces ge-0/0/1 unit 100 vlan-id 100
user@vmx# set interfaces ge-0/0/1 unit 100 family inet address 10.1.1.1/24

```

3. Configure the classifiers.

```

[edit]
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class voice1 loss-priority

```

low code-points 000

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class video1 loss-
priority low code-points 001
```

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class data1 loss-priority
low code-points 010
```

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class data2 loss-priority
```

low code-points 011

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class data3 loss-priority
```

low code-points 100

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class data4 loss-priority
```

low code-points 101

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class voice2 loss-priority
```

low code-points 110

```
user@vmx# set class-of-service classifiers inet-precedence vlan_tos forwarding-class video2 loss-
priority low code-points 111
```

4. Configure the traffic control profiles.

```
[edit]
user@vmx# set class-of-service traffic-control-profiles ge_0_0_1_vlan_100_tcp shaping-rate 50m
user@vmx# set class-of-service traffic-control-profiles ge_0_0_1_vlan_100_tcp scheduler-map
vlan_smap
```

5. Map the traffic control profiles to their respective interface.

```
[edit]
user@vmx# set class-of-service interfaces ge-0/0/1 unit 100 output-traffic-control-profile
ge_0_0_1_vlan_100_tcp
user@vmx# set class-of-service interfaces ge-0/0/0 unit 100 classifiers inet-precedence vlan_tos
```

6. Configure the scheduler maps.

```
[edit]
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class voice1 scheduler
sched_voice1
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class video1 scheduler
sched_video1
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class data1 scheduler
```

```

sched_data1
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class data2 scheduler
sched_data2
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class data3 scheduler
sched_data3
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class data4 scheduler
sched_data4
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class voice2 scheduler
sched_voice2
user@vmx# set class-of-service scheduler-maps vlan_smap forwarding-class video2 scheduler
sched_video2

```

7. Configure the schedulers.

```

[edit]
user@vmx# set class-of-service schedulers sched_voice1 transmit-rate 15m
user@vmx# set class-of-service schedulers sched_video1 transmit-rate 15m
user@vmx# set class-of-service schedulers sched_data1 transmit-rate 5m
user@vmx# set class-of-service schedulers sched_data2 transmit-rate 5m
user@vmx# set class-of-service schedulers sched_data3 transmit-rate 5m
user@vmx# set class-of-service schedulers sched_data4 transmit-rate 5m
user@vmx# set class-of-service schedulers sched_voice2 transmit-rate 10m
user@vmx# set class-of-service schedulers sched_video2 transmit-rate 10m

```

RELATED DOCUMENTATION

[Configuring Hierarchical CoS on vMX | 56](#)

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)

[Configuring Hierarchical CoS on vMX | 56](#)

Configuring Four-Level Hierarchical Scheduling on vMX

Starting with Junos OS Release 17.3R1, four-level hierarchical scheduling for up to 16 level 2 CoS scheduler nodes is supported on vMX routers. The level 2 node maps to the interface set or VLAN (logical interface). Two of the level 2 nodes are used for control traffic. If you configure less than four nodes, no commit errors are displayed but there are not enough nodes for other applications to use. Different interfaces can have a different number of level 2 nodes. The interface can be an inline service interface.

To configure four-level hierarchical scheduling:

1. Hierarchical CoS is disabled by default. Configure flexible queuing to enable CoS.

```
[edit]
user@vmx# set chassis fpc 0 flexible-queuing-mode
```

NOTE: The FPC reboots if you enable flexible queuing.

2. Enable hierarchical scheduling.

```
[edit]
user@vmx# set interfaces interface-name implicit-hierarchy
```

3. Set the maximum number of hierarchical scheduling levels for node scaling to 3. If the **maximum-hierarchy-levels** option is not configured, it is automatically set to 2.

```
[edit]
user@vmx# set interfaces interface-name hierarchical-scheduler maximum-hierarchy-levels 3
```

4. Specify the maximum number of level 2 scheduler nodes; only 1, 2, 4, 8, and 16 are valid values. The default value is 4. We recommend that you do not configure less than four nodes because two of the nodes are used for control traffic.

```
[edit]
user@vmx# set interfaces interface-name hierarchical-scheduler maximum-l2-nodes number-of-nodes
```

For example:

```
[edit]
user@vmx# set interfaces ge-0/0/0 hierarchical-scheduler maximum-l2-nodes 4
```

NOTE: This configuration must be present before you reboot the FPC.

RELATED DOCUMENTATION

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)

[Configuring Hierarchical CoS on vMX | 56](#)

Packet Loss Priority and Drop Profiles on vMX

IN THIS SECTION

- [Limitations | 67](#)

vMX handles packet priorities within a queue by assigning a threshold to each loss priority within a queue and dropping new packets of that loss priority level when the queue depth exceeds the threshold. When the queue becomes oversubscribed, packets of lower priority are dropped to ensure that there is room in the queue for packets of higher priority.

Packet loss priority has four loss priority levels:

- low
- medium-low
- medium-high
- high

vMX supports three thresholds so the medium-low and medium-high loss priority levels are grouped together. vMX maps the packet loss priority to tricolor marking as follows:

Packet Loss Priority	Color
low	green
medium-low	yellow
medium-high	yellow
high	red

vMX drop profiles define the threshold within a queue for a given loss priority as the fill level value associated with the drop probability of 100 percent. If you do not specify a drop probability of 100 percent in the drop profile, the threshold defaults to 100 percent. All other fill level values are ignored. These drop profiles can be referenced by the scheduler to evaluate packets with different loss priority settings.

You can set packet loss priority for packets using behavior aggregate (BA) classifiers, firewall filters, or firewall policers.

Limitations

vMX has the following limitations for supporting drop profiles and packet loss priority:

- If you do not apply drop profiles to the queue, then packets are tail dropped.
- The **show interface queue** command does not display separate drop rates for the medium-high PLP and medium-low PLP because they both map to yellow. All yellow drop rates appear as medium-high drops.

RELATED DOCUMENTATION

[Managing Congestion Using Drop Profiles and Packet Loss Priorities on vMX | 68](#)

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)

Managing Congestion Using Drop Profiles and Packet Loss Priorities on vMX

IN THIS SECTION

- [Configuring Drop Profiles | 68](#)
- [Configuring Schedulers with Drop Profiles | 69](#)

When you are configuring CoS, you can manage congestion by configuring drop profiles to specify the thresholds for packet loss priority. You reference the drop profiles in the scheduler configuration to assign a drop profile to the loss priority setting.

To configure how packet loss priority is handled for queues, perform these tasks:

Configuring Drop Profiles

Drop profiles specify the threshold for a given loss priority.

NOTE: The threshold for the loss priority assigned this drop profile is the **fill-level** value associated with the **drop-probability** of 100. If you do not specify a drop probability of 100 percent in the drop profile, the fill level defaults to 100 percent. All other fill levels are ignored.

To specify the drop profile, include the **drop-profiles** statement at the **[edit class-of-service]** hierarchy level.

```
[edit]
user@vmx# set class-of-service drop-profiles profile-name
```

To specify the threshold for the loss priority, include the **fill-level** and **drop-probability** statements at the **[edit class-of-service drop-profiles *profile-name*]** hierarchy level.

```
[edit class-of-service drop-profiles profile-name]
user@vmx# set fill-level percentage drop-probability percentage
```

For example, the **dpLow** drop profile specifies a threshold of 100 percent, the **dpMed** drop profile specifies a threshold of 75 percent, and the **dpHigh** drop profile specifies a threshold of 50 percent.

```
[edit]
user@vmx# set class-of-service drop-profiles dpLow fill-level 100 drop-probability 100
user@vmx# set class-of-service drop-profiles dpMed fill-level 75 drop-probability 100
user@vmx# set class-of-service drop-profiles dpHigh fill-level 50 drop-probability 100
```

Configuring Schedulers with Drop Profiles

The drop profile map contains the mapping of loss priority and protocol type to configured drop profiles. You can associate multiple drop profile maps with a scheduler.

NOTE: If you do not apply drop profiles to the queue, then packets are tail dropped.

To specify the drop profile map, include the **drop-profile-map** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.

```
[edit class-of-service schedulers scheduler-name]
user@vmx# set drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
protocol any drop-profile profile-name
```

For example, the **sched-be** scheduler applies the **dpLow** drop profile to packets with low loss priority for any protocol type, applies the **dpMed** drop profile to packets with medium-high loss priority for any protocol type, and applies the **dpHigh** drop profile to packets with high loss priority for any protocol type.

```
[edit class-of-service schedulers sched-be]
user@vmx# set drop-profile-map loss-priority low protocol any drop-profile dpLow
user@vmx# set drop-profile-map loss-priority medium-high protocol any drop-profile dpMed
user@vmx# set drop-profile-map loss-priority high protocol any drop-profile dpHigh
```

RELATED DOCUMENTATION

[Packet Loss Priority and Drop Profiles on vMX | 66](#)

[CoS on vMX Overview | 52](#)

[CoS Features and Limitations on vMX | 54](#)



CHAPTER

Troubleshooting vMX

[Viewing VFP Statistics](#) | 72

[Viewing VFP Log Files](#) | 73

[Troubleshooting VFP and VCP Connection Establishment](#) | 74

Viewing VFP Statistics

You can view the VFP statistics from a Web browser. The displayed statistics are not absolute counters; they are relative to the start of the HTTP session and start as all zero counters.

The RPIO Stats and Hostif Stats sections provide statistics about the internal communication between the VCP and the VFP. The RPIO session uses ports 3000 and 3001 and the HostIF session uses port 3002.

The Port Stats section provides statistics about the packets received from and transmitted to the NIC interfaces.

- There is a receive (**rx**) and transmit (**tx**) line for each port. Port 0 maps to the ge-0/0/0 interface, port 1 maps to the ge-0/0/1 interface, and so forth. rx0 displays statistics for packets received from port 0 and tx1 displays statistics for packets transmitted to port 1.
- Errors are miscellaneous errors reported by the physical layer NIC.

The Ring Stats section provides statistics about packet processing.

- There is an I/O thread (**io**) for packets received from a port.
- There is a Worker thread (**wk**) for each CPU core.
- The host interface (**host**) sends protocol packets to the VCP.
- The queue processes the packets. The columns provide this information about the queues:
 - The Producer and Consumer columns display the source and destination that generate packets for this queue. The values can be **io**, **wk**, **tx**, or **host**.
 - The Priority column displays the priority of the queue. The values can be **Normal** or **High** (only for control packets).
 - The Free and Used columns display the queue occupancy. The queue has 1024 entries.
 - The Enqueues and Dequeues columns display the number of queue operations.
 - The Drops column indicates whether the queue is being drained fast enough.

To view the statistics:

1. By default, you cannot log in to the Web browser window without configuring the username and password credentials and enabling HTTP access.

From the VFP console, configure the username and password by invoking the `/home/pfe/riot/vfp_util.sh -setpass` command.

```
root@vfp-vmx1:/home/pfe/riot# ./vfp_util.sh -setpass
Enter new Username: pfe
Enter new Password:
Re-enter Password:
Password successfully changed
root@vfp-vmx1:/home/pfe/riot#
```

To enable HTTP access, invoke this command.

```
root@vfp-vmx1:/home/pfe/riot# ./vfp_util.sh -http_enable
```

2. Navigate to `http://vfp-mgmt-ip:8080/`, where *vfp-mgmt-ip* is the management IP address for the VFP VM.
3. When prompted, enter **pfe** as the username and the password configured in Step 1.
4. View the statistics displayed in the browser window.
5. After troubleshooting, you can disable HTTP access to improve security with this command:

```
root@vfp-vmx1:/home/pfe/riot# ./vfp_util.sh -http_disable
```

RELATED DOCUMENTATION

[Viewing VFP Log Files | 73](#)

[vMX Overview | 2](#)

Viewing VFP Log Files

The VFP saves the following files:

- VFP log files are saved in the `/var/log` directory.
- VFP crash files are automatically saved in the VCP `/var/crash` directory.

To view the VFP log or crash files:

1. Log in to the VFP console by using the `telnet esxi-server-ip-address port-number` command, where *port-number* is the port number specified for the Port URI in the serial port connection configuration for the VFP VM.
2. Navigate to the appropriate directory to determine whether there are any files to view.

```
# cd /var/crash
# ls -l
```

```
-rwxr-xr-x 1 root root 864678 Jan  4 02:14 core.riot.1420366466.8271.gz
```

3. (Optional) If necessary, unzip the file and view it using GDB.

```
# gunzip core.riot.1420366466.8271.gz
# gdb core.riot.1420366466.8271
```

RELATED DOCUMENTATION

[Viewing VFP Statistics | 72](#)

[vMX Overview | 2](#)

Troubleshooting VFP and VCP Connection Establishment

IN THIS SECTION

 [Purpose | 75](#)

Purpose

When the VCP and VFP connection is established, the **show interfaces terse** command in the VCP CLI displays the ge-0/0/x interfaces and the following messages appear in the VFP syslog file:

```
RPIO: Accepted connection from 128.0.0.1:50896 <-> vPFE:3000
RPIO: Accepted connection from 128.0.0.1:56098 <-> vPFE:3000
HOSTIF: Accepted connection
```

If the VCP cannot connect to the VFP, the VFP syslog file does not display the **RPIO** and **HOSTIF** messages.

Action

Run the **request chassis fpc slot 0 restart** command from the VCP CLI to restart the FPC. If an **FPC is in transition** error message is displayed, then run **restart chassis-control**.

If these commands do not correct the problem, verify whether the VCP can ping the VFP from the **routing-instance __juniper_private1__**. The three management interfaces (for the host, VCP VM, and VFP VM) connected to the internal bridge should be able to reach each other. For example:

```
root> ping 128.0.0.16 routing-instance __juniper_private1__
PING 128.0.0.16 (128.0.0.16): 56 data bytes
64 bytes from 128.0.0.16: icmp_seq=0 ttl=64 time=0.273 ms
64 bytes from 128.0.0.16: icmp_seq=1 ttl=64 time=0.606 ms
```

If the VCP cannot ping the VFP, perform these tasks:

1. Make sure that both the VCP and VFP have the network adapter connected to the internal bridge vSwitch (br-int).
2. Verify that the MAC address of em1 on the VCP matches the MAC address shown in the Virtual Machine Properties.

To determine the MAC address of the em1 interface:

```
root@% ifconfig em1 | grep "curr media" | awk '{print $NF}'
```

3. Verify that the MAC address of the interface on the VFP matches the MAC address shown in the Virtual Machine Properties.

To determine the MAC address of the interface:

```
root@localhost:/var/log# ifconfig int | grep "HWaddr" | awk '{print $NF}'
```

If the problem persists, contact the Juniper Networks Technical Assistance Center (JTAC).