



Cisco UCS Manager VM-FEX for KVM GUI Configuration Guide, Release 3.1

First Published: 2016-01-20

Last Modified: 2016-09-17

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface v

Audience v

Conventions v

Related Documentation vii

Obtaining Documentation and Submitting a Service Request vii

CHAPTER 1

Introduction 1

Overview of Virtualization 1

Overview of Cisco Virtual Machine Fabric Extender 1

Virtualization with a Virtual Interface Card Adapter 2

Single Root I/O Virtualization 2

VM-FEX for KVM 3

Overview of VM-FEX for KVM 3

Cisco UCS Manager Components 3

KVM Components 4

Driver Topologies 5

CHAPTER 2

Configuring VM-FEX for KVM 9

Guidelines and Prerequisites for KVM 9

Configuring VM-FEX for SR-IOV with MacVTap Topology 10

Configuring VM-FEX for SR-IOV Passthrough Topology 11

Configuring the VM Interface 11

Activating Intel VT-d in the Kernel 15

CHAPTER 3

Configuring a Service Profile with VM-FEX 17

Configuring Dynamic vNIC Connection Policies 17

Dynamic vNIC Connection Policy 17

Creating a Dynamic vNIC Connection Policy	18
Changing a Dynamic vNIC Connection Policy	19
Deleting a Dynamic vNIC Connection Policy	20
Viewing Dynamic vNIC Properties in a VM	20

CHAPTER 4**Configuring Port Profiles 23**

Port Profiles	23
Creating a Port Profile	23
Modifying the VLANs in a Port Profile	25
Changing the Native VLAN for a Port Profile	26
Adding a VLAN to a Port Profile	26
Deleting a VLAN from a Port Profile	26
Deleting a Port Profile	27
Port Profile Client	27
Creating a Profile Client	27
Modifying a Profile Client	28
Deleting a Profile Client	29



Preface

- [Audience, page v](#)
- [Conventions, page v](#)
- [Related Documentation, page vii](#)
- [Obtaining Documentation and Submitting a Service Request, page vii](#)

Audience

This guide is intended primarily for data center administrators with responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in this font . Main titles such as window, dialog box, and wizard titles appear in this font .
Document titles	Document titles appear in <i>this font</i> .
TUI elements	In a Text-based User Interface, text the system displays appears in <i>this font</i> .
System output	Terminal sessions and information that the system displays appear in <i>this font</i> .

Text Type	Indication
CLI commands	CLI command keywords appear in this font . Variables in a CLI command appear in <i>this font</i> .
[]	Elements in square brackets are optional.
{x y z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
< >	Nonprinting characters such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

**Note**

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the document.

**Tip**

Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

**Timesaver**

Means *the described action saves time*. You can save time by performing the action described in the paragraph.

**Caution**

Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

**Warning****IMPORTANT SAFETY INSTRUCTIONS**

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

UCS Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: <http://www.cisco.com/go/unifiedcomputing/c-series-doc>.

Other Documentation Resources

An ISO file containing all B and C-Series documents is available at the following URL: <http://www.cisco.com/cisco/software/type.html?mdfid=283853163&flowid=25821>. From this page, click **Unified Computing System (UCS) Documentation Roadmap Bundle**.

The ISO file is updated after every major documentation release.

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly [What's New in Cisco Product Documentation](#), which also lists all new and revised Cisco technical documentation.

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.



Introduction

- [Overview of Virtualization, page 1](#)
- [Overview of Cisco Virtual Machine Fabric Extender, page 1](#)
- [Virtualization with a Virtual Interface Card Adapter, page 2](#)
- [Single Root I/O Virtualization, page 2](#)
- [VM-FEX for KVM, page 3](#)

Overview of Virtualization

Virtualization allows you to create multiple Virtual Machines (VMs) to run in isolation, side by side on the same physical machine.

Each virtual machine has its own set of virtual hardware (RAM, CPU, NIC) upon which an operating system and fully configured applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components.

In a virtual machine, both hardware and software are encapsulated in a single file for rapid provisioning and moving between physical servers. You can move a virtual machine, within seconds, from one physical server to another for zero-downtime maintenance and continuous workload consolidation.

The virtual hardware makes it possible for many servers, each running in an independent virtual machine, to run on a single physical server. The advantages of virtualization include better use of computing resources, greater server density, and seamless server migration.

Overview of Cisco Virtual Machine Fabric Extender

A virtualized server implementation consists of one or more VMs that run as guests on a single physical server. The guest VMs are hosted and managed by a software layer called the hypervisor or virtual machine manager (VMM). Typically, the hypervisor presents a virtual network interface to each VM and performs Layer 2 switching of traffic from a VM to other local VMs or to another interface to the external network.

Working with a Cisco virtual interface card (VIC) adapter, the Cisco Virtual Machine Fabric Extender (VM-FEX) bypasses software-based switching of VM traffic by the hypervisor for external hardware-based

switching in the fabric interconnect. This method reduces the load on the server CPU, provides faster switching, and enables you to apply a rich set of network management features to local and remote traffic.

VM-FEX extends the IEEE 802.1Qbh port extender architecture to the VMs by providing each VM interface with a virtual Peripheral Component Interconnect Express (PCIe) device and a virtual port on a switch. This solution allows precise rate limiting and quality of service (QoS) guarantees on the VM interface.

Virtualization with a Virtual Interface Card Adapter

A Cisco VIC adapter is a converged network adapter (CNA) that is designed for both bare metal and VM-based deployments. The VIC adapter supports static or dynamic virtualized interfaces, which includes up to 128 virtual network interface cards (vNICs).

There are two types of vNICs used with the VIC adapter—static and dynamic. A static vNIC is a device that is visible to the OS or hypervisor. Dynamic vNICs are used for VM-FEX by which a VM is connected to a veth port on the Fabric Interconnect.

VIC adapters support VM-FEX to provide hardware-based switching of traffic to and from virtual machine interfaces.

Single Root I/O Virtualization

Single Root I/O Virtualization (SR-IOV) allows multiple VMs running a variety of guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the network adapter, bypassing the hypervisor for increased network throughput and lower server CPU burden. Recent x86 server processors include chipset enhancements, such as Intel VT-x technology, that facilitate direct memory transfers and other operations required by SR-IOV.

The SR-IOV specification defines two device types:

- Physical Function (PF)—Essentially a static vNIC, a PF is a full PCIe device that includes SR-IOV capabilities. PFs are discovered, managed, and configured as normal PCIe devices. A single PF can provide management and configuration for a set of virtual functions (VFs).
- Virtual Function (VF)—Similar to a dynamic vNIC, a VF is a full or lightweight virtual PCIe device that provides at least the necessary resources for data movements. A VF is not managed directly but is derived from and managed through a PF. One or more VFs can be assigned to a VM.

SR-IOV is defined and maintained by the Peripheral Component Interconnect Special Interest Group (PCI-SIG), an industry organization that is chartered to develop and manage the PCI standard. For more information about SR-IOV, see the following URL:

<http://www.intel.com/content/www/us/en/pci-express/pci-sig-sr-iov-primer-sr-iov-technology-paper.html>

Hypervisors that support SR-IOV include Linux KVM and Microsoft Hyper-V.

The following Cisco Virtual Interface Cards support SR-IOV with VM-FEX:

- Cisco UCS M81KR Virtual Interface Card
- Cisco UCS P81E Virtual Interface Card
- Cisco UCS Virtual Interface Card 1280
- Cisco UCS Virtual Interface Card 1240

- Cisco UCS Virtual Interface Card 1225
- Cisco UCS Virtual Interface Card 1225T
- Cisco UCS Virtual Interface Card 1227
- Cisco UCS Virtual Interface Card 1227T
- Cisco UCS Virtual Interface Card 1340
- Cisco UCS Virtual Interface Card 1380
- Cisco UCS Virtual Interface Card 1385

VM-FEX for KVM

Overview of VM-FEX for KVM

The Kernel-based Virtual Machine (KVM) is a virtualization package for Linux on an x86 hardware platform. KVM uses x86 hardware virtualization extensions (for example, Intel VT-*x*) to implement a hypervisor that hosts VMs as userspace processes. Cisco UCS servers support the KVM-based Red Hat Enterprise Virtualization (RHEV) as the hypervisor in a server virtualization system.

With VM-FEX for KVM, the RHEV hypervisor performs no switching of VM traffic. Working with an installed VIC adapter, the hypervisor acts as an interface virtualizer and performs the following functions:

- For traffic going from a VM to the VIC, the interface virtualizer identifies the source vNIC so that the VIC can explicitly tag each packet that is generated by that vNIC.
- For traffic that is received from the VIC, the interface virtualizer directs the packet to the specified vNIC.

All switching is performed by the external fabric interconnect, which can switch not only between physical ports, but also between virtual interfaces (VIFs) that correspond to the vNICs on the VMs.

For more information about KVM, see the following URL: <http://www.linux-kvm.org>.

Cisco UCS Manager Components

Cluster

The Cisco UCS cluster is a grouping of hypervisors that can be distributed across multiple hosts. In a KVM system, the cluster is analogous to the distributed virtual switch (DVS) in a VMware ESX system.

In the current Cisco UCS KVM implementation, the cluster defines the scope of the port profile and is the boundary of the migration domain. When multiple KVM hosts are associated to a cluster, you can migrate a VM from one host to another within the cluster.



Note

In the current Cisco UCS implementation of VM-FEX for KVM, only one cluster, the default cluster, is used. Although you can create additional clusters, you can specify only the default cluster for a VM on the KVM host.

Port Profiles

Port profiles contain the properties and settings that are used to configure virtual interfaces in Cisco UCS. The port profiles are created and administered in Cisco UCS Manager.



Important

After a port profile is created, assigned to, and actively used by a cluster, any changes made to the networking properties of the port profile in Cisco UCS Manager are immediately applied to the cluster with no need for a host reboot.

Port Profile Client

The port profile client is a cluster to which a port profile is applied.



Note

In the current Cisco UCS implementation of VM-FEX for KVM, the default cluster is the only available port profile client.

KVM Components

Hypervisor

The hypervisor supports multiple VMs that run a variety of guest operating systems by providing connectivity between the VMs and the network. The hypervisor for KVM is a host server with Red Hat Enterprise Linux (RHEL) installed. The earliest supported release for VM-FEX is RHEL 6.1, but some features (such as SR-IOV) require a later version.

The hypervisor must have a Cisco VIC adapter installed.

For more information about virtualization using Red Hat Enterprise Linux, see the *Red Hat Enterprise Virtualization for Servers Installation Guide* available at the following URL: <http://www.redhat.com>.

libvirt

Libvirt is an open source toolkit that allows you to manage various virtualization technologies such as KVM, Xen, and VMware ESX. Libvirt, which runs on the hypervisor as a service named libvirtd, provides a command-line interface (virsh) and provides the toolkit for a graphical user interface package (virt-manager).

Each virtual machine created and managed by libvirt is represented in the form of a domain XML file.

For more information about the libvirt virtualization API, see the following URL: <http://www.libvirt.org>.

For more information about the virsh CLI, see the following URLs:

- <http://linux.die.net/man/1/virsh>
- <http://www.libvirt.org/virshcmdref.html>

MacVTap

MacVTap is a Linux driver that allows the direct attachment of a VM's vNIC to a physical NIC on the host server.

For more information about the MacVTap driver, see the following URL: <http://virt.kernelnewbies.org/MacVTap>.

VirtIO

The VirtIO paravirtualized network driver (virtio-net) runs in the guest operating system of the VM and provides a virtualization-aware emulated network interface to the VM.

For more information about the VirtIO driver, see the following URL: <http://wiki.libvirt.org/page/Virtio>.

Driver Topologies

Several driver topologies (modes) are available to implement a VM-FEX connection between a VM vNIC and the host VIC adapter. In each of these topologies, VM traffic is sent only to or from the VIC adapter. Traffic from one VM to another VM on the same host must first exit the host for switching by the external fabric interconnect.



Note

In any topology, the configuration of the Quick EMUlator (QEMU) PCI layer might limit the number of PCI devices that the host can assign to a VM.

MacVTap Direct (Private)

The MacVTap Linux driver is installed in the hypervisor (VMM) and connects each VM's VirtIO interface to a physical PCIe port of the VIC adapter. The MacVTap driver mode is private, which means that all VM traffic is sent directly to and from the host adapter with external switching. The number of supported VMs is limited to the number of VIC adapter ports. Live migration is supported.



Note

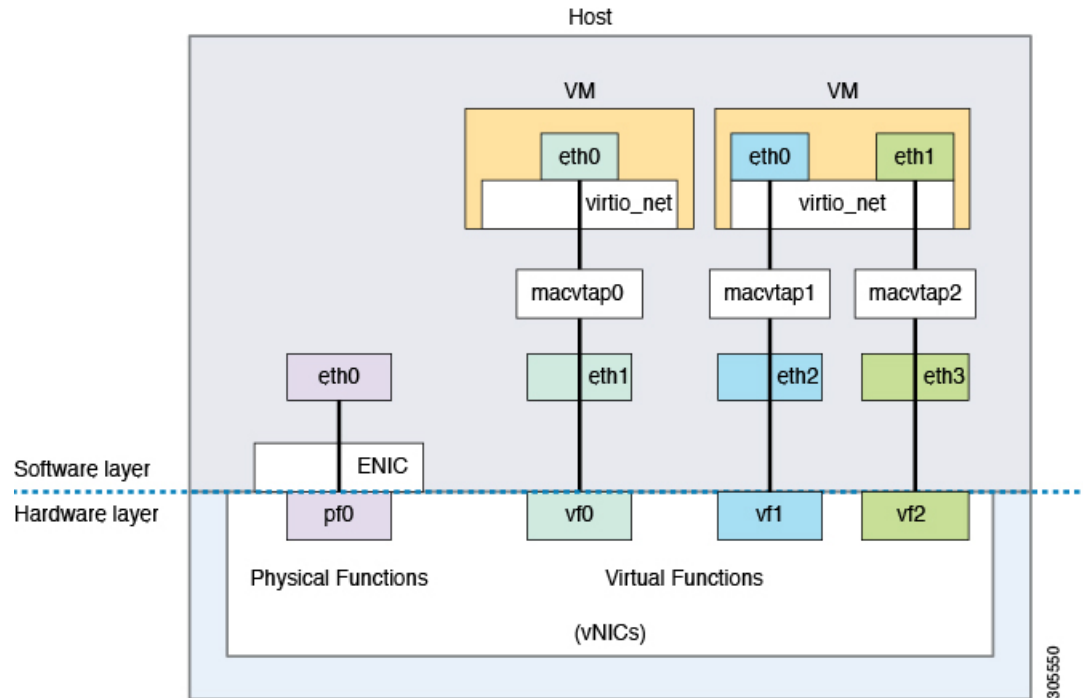
Beginning with Cisco UCS Release 2.1, the MacVTap Direct (Private) topology is no longer supported.

SR-IOV with MacVTap Passthrough (Emulation Mode)

The MacVTap Linux driver is installed in the hypervisor and connects each VM's VirtIO interface to a VF on an SR-IOV-capable VIC adapter. The MacVTap driver mode is 'passthrough' and all VM traffic is sent to and from the VF. When we apply a port profile to a VF, libvirt determines the PF associated with the VF, and it configures the VF going through the PF. This topology is also known as MacVTap passthrough (emulation

mode). An example of SR-IOV with MacVTap passthrough is shown in Figure 1. Figure 1 is a simplified version of the hardware and software components.

Figure 1: Figure 1



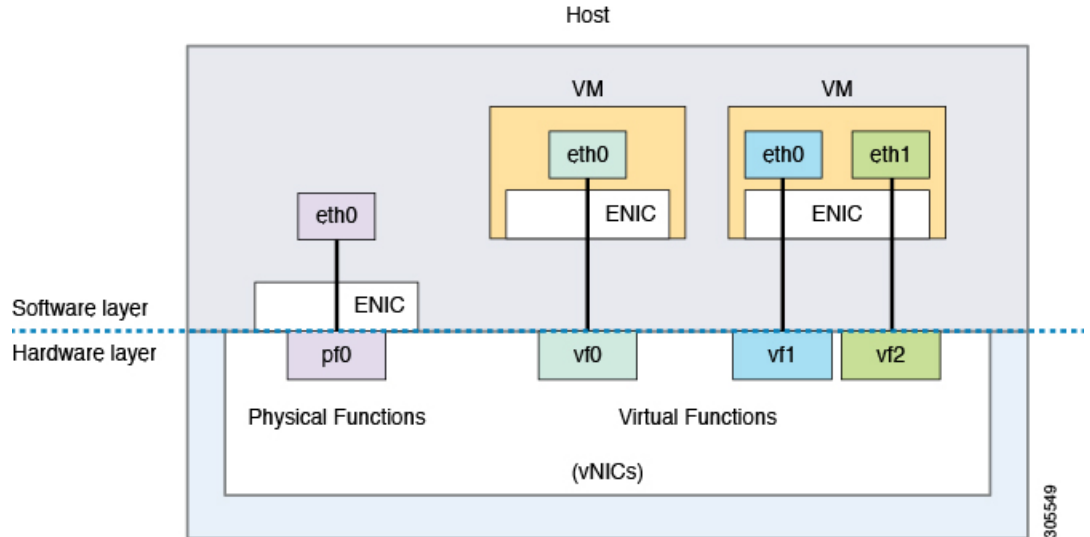
The maximum number of supported VMs is determined by the number of VFs provided by the VIC adapter. The number of VFs that you can assign to a PF might be further limited by the host Netlink protocol implementation (the limit is typically between 22 and 32 VFs per PF, depending on the OS version). Live migration is supported.

SR-IOV VF Passthrough (Hostdev Mode)

The MacVTap and VirtIO drivers are not used. Instead, the Ethernet driver (`enic`) of the VIC adapter is installed in the VM kernel and connects directly to a VF. You can configure the VF through the associated PF using `libvirt`. In `libvirt` documentation, this topology is called `hostdev` mode. This topology is also known as PCI passthrough. The number of supported VMs is determined by the number of VFs provided by the VIC adapter.

Live migration is not supported. An example of SR-IOV with VF passthrough is shown in Figure 2. Figure 2 is a simplified version of the hardware and software components.

Figure 2: Figure 2





CHAPTER 2

Configuring VM-FEX for KVM

- [Guidelines and Prerequisites for KVM, page 9](#)
- [Configuring VM-FEX for SR-IOV with MacVTap Topology, page 10](#)
- [Configuring VM-FEX for SR-IOV Passthrough Topology, page 11](#)
- [Configuring the VM Interface, page 11](#)
- [Activating Intel VT-d in the Kernel, page 15](#)

Guidelines and Prerequisites for KVM

Consider the following guidelines and prerequisites when configuring Kernel-based Virtual Machine (KVM):

- The host must be managed by Cisco UCS Manager Release 2.1 or later.
- On Red Hat Enterprise Linux (RHEL) hosts, disable generic receive offload (GRO) by using the `ethtool-K interface gro off` command. This issue occurs because Microsoft Windows VIRTIO does not support GRO, which results in very poor Ethernet performance compared with Linux VMs.
- The host operating system must be RHEL with KVM support.
 - The Single Root I/O Virtualization (SR-IOV) with MacVTap topology requires RHEL 6.2 or later.
 - The SR-IOV passthrough topology requires RHEL 6.3 or later.

For more information about installing RHEL with KVM, see the *Red Hat Enterprise Virtualization for Servers Installation Guide*.

- The host must have libvirt with virsh or virt-manager installed for creating and managing the VMs.
- One or more Cisco VIC adapters must be installed in the host.

For more information about installing a Cisco VIC adapter, see the *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Consider the following guidelines and prerequisites when configuring an SR-IOV topology:

- Intel VT-x and VT-d processor extensions for virtualization must be enabled in the host BIOS.

For more information about configuring Cisco UCS server BIOS settings, see the *Cisco UCS Manager GUI Configuration Guide*.

- For SR-IOV topologies, configure a dynamic connection policy in a service profile. Apply the service profile on a static vNIC to specify the number of VFs, the fabric preference, and the adapter policy. The static vNIC becomes a PF when you configure one or more VFs under it. VFs are provisioned as dynamic vNICs.
- All VF-based dynamic vNICs must be provisioned on the same physical adapter as the parent static vNIC (PF).
- When you upgrade Cisco UCS Manager to an SR-IOV capable release, the existing static and dynamic vNICs are not automatically enabled for SR-IOV. To convert to SR-IOV, you must disable any existing dynamic connection policy in the service profile and then specify a reference to a dynamic connection policy under a static vNIC.

Configuring VM-FEX for SR-IOV with MacVTap Topology

Before You Begin

Prepare the host server as described in [Guidelines and Prerequisites for KVM](#), on page 9.

Procedure

-
- Step 1** In Cisco UCS Manager, configure a service profile for VM-FEX for KVM. Create or modify a dynamic vNIC connection policy.
- For more information, see [Configuring a Service Profile with VM-FEX](#), on page 17.
- Step 2** In Cisco UCS Manager, define a port profile and associate it with a port profile client. Create a port profile to define the properties and settings used to configure the virtual interfaces. For KVM, you must select the default cluster as the port profile client.
- For more information, see [Configuring Port Profiles](#), on page 23.
- Step 3** On each KVM server, use `virsh` or `virt-manager` to create one or more virtual machines (VMs). For more information about installing VMs using these libvirt-based utilities, see the documents listed in [KVM Components](#), on page 4.
- Note** When creating a VM using `virsh`, or when editing the VM domain XML descriptor file, use care when entering data such as a universally unique identifier (UUID), as you will receive no indication of incorrect data values or formats.
- Step 4** For each VM, edit the domain XML descriptor file (and any network XML files, if present) to configure a vNIC interface that is directly attached to the VIC and uses the port profile defined in Cisco UCS Manager. For more information about configuring a VM interface, see [Configuring the VM Interface](#), on page 11.
- Step 5** On each VM, install the VirtIO paravirtualized network driver (`virtio-net`) for the guest operating system. Recent versions of most common operating systems provide default `virtio-net` drivers. For more information, contact Red Hat or the provider of the guest operating system.
-

Configuring VM-FEX for SR-IOV Passthrough Topology

Before You Begin

Prepare the host server as described in [Guidelines and Prerequisites for KVM](#), on page 9.

Procedure

- Step 1** In Cisco UCS Manager, configure a service profile for VM-FEX for KVM. Create or modify a dynamic vNIC connection policy. For more information, see [Configuring a Service Profile with VM-FEX](#), on page 17.
- Step 2** In Cisco UCS Manager, define a port profile and associate it with a port profile client. Create a port profile to define the properties and settings used to configure the virtual interfaces. For KVM, you must select the default cluster as the port profile client. For more information, see [Configuring Port Profiles](#), on page 23.
- Step 3** On each KVM server, use virsh or virt-manager to create one or more virtual machines (VMs). For more information about installing VMs using these libvirt-based utilities, see the documents listed in [KVM Components](#), on page 4.
- Note** When creating a VM using virsh, or when editing the VM domain XML descriptor file, use care when entering data such as a universally unique identifier (UUID), as you will receive no indication of incorrect data values or formats.
- Step 4** On each VM, install an enic driver that supports an SR-IOV VF. With RHEL 6.3 or later, use the inbox enic driver.
- Step 5** For each VM, edit the domain XML descriptor file (and any network XML files, if present) to configure a vNIC interface that is directly attached to the VIC and uses the port profile defined in Cisco UCS Manager. For more information about configuring a VM interface, see [Configuring the VM Interface](#), on page 11.
- Step 6** On the KVM host, activate the Intel VT-d extensions. For more information about activating the VT-d extensions, see [Activating Intel VT-d in the Kernel](#), on page 15.
-

Configuring the VM Interface

After creating a VM using a libvirt-based utility, you can add the network configuration, including the port profile information, either to the domain XML of the VM or to a separate network XML file that you can reference from the domain XML. One of the advantages of adding the configuration information to a separate network XML file is that you can specify a pool of devices. For more information about the network XML file components and attributes, see <http://libvirt.org/formatnetwork.html>.

For more information about the domain XML file components and attributes, see the libvirt documentation at <http://libvirt.org/formatdomain.html#elementsNICs>.

Procedure

-
- Step 1** Shut down the VM to be configured.
- Step 2** Using the virsh editor, open the domain XML file of the VM for editing.

Example:

This example opens a domain XML file for editing in the virsh editor:

```
[root@chassis1blade5 qemu]# virsh edit vml-rhel6.2
```

- Step 3** In the devices section of the domain XML file, add an interface element that describes a vNIC for the VM. The components and attributes of the interface element are described in the Example section.
- Step 4** Restart the VM.
-

Example for SR-IOV with MacVTap Mode

This example shows an interface element added to the domain XML file of a VM for connection in SR-IOV with MacVTap (MacVTap Passthrough) topology:

```
<domain type='kvm'>
  <name>vml-rhel6.2</name>
  ...
  <devices>
    ...
    <interface type='direct'>
      <mac address='01:23:45:67:89:ab' />
      <source dev='eth4' mode='passthrough' />
      <virtualport type='802.1Qbh'>
        <parameters profileid='my-port-profile-3' />
      </virtualport>
      <model type='virtio' />
      <driver name='vhost' />
    </interface>
    ...
  </devices>
  ...
</domain>
```

This list describes the components and attributes of the interface element:

- `interface type='direct'`

The `direct` type attribute value selects a direct logical attachment of the vNIC to the physical interface of the hypervisor, using the MacVTap driver.

- `mac address='01:23:45:67:89:ab'`

Explicit specification of the MAC address is optional. Enter a MAC address obtained from your network administrator. If this line is omitted, libvirt generates a MAC address for the vNIC.



Note We recommend that you do not assign a MAC address used by another VM, even if that VM is currently shut down or is no longer used. If you must reuse a MAC address from a previous VM, make sure that the retention timer has expired and ensure that the previous VM is no longer present in the Cisco UCS Manager view.

- `source dev='eth4' mode='passthrough'`

The `passthrough` mode attribute value specifies that each VM is connected to the network by a macvtap direct connection with a virtual function (VF). The source interface must be a VF, and not a physical function (PF).

- `virtualport type='802.1Qbh'`

The `802.1Qbh` type attribute value specifies that the vNIC is connected to an 802.1Qbh extended port for external switching.

- `parameters profileid='my-port-profile-3'`

This line specifies the name of the port profile to be associated with the interface. The port profile name is case sensitive. The specified port profile must be already defined in Cisco UCS Manager and use the naming syntax described in [Creating a Port Profile, on page 23](#).

- `model type='virtio'`

This line specifies that the guest interface uses the VirtIO paravirtualized front-end device driver.

- `driver name='vhost'`

This line specifies that, for higher performance, the host side interface uses the vhost kernel back-end device driver and not the qemu userspace back-end driver.

Example for SR-IOV Passthrough Mode

This example shows an interface element that is added to the domain XML file of a VM for a connection in SR-IOV Passthrough topology:

```
<domain type='kvm'>
  <name>vml-rhel6.3</name>
  ...
  <devices>
    ...
    <interface type='hostdev' managed='yes'>
      <source>
        <address type='pci' domain='0' bus='0x09' slot='0x0' function='0x01' />
      </source>
      <mac address='01:23:45:67:89:ab' />
      <virtualport type='802.1Qbh'>
        <parameters profileid='my-port-profile-3' />
      </virtualport>
    </interface>
    ...
  </devices>
  ...
</domain>
```

This list describes the components and attributes of the interface element that differ from those described in the SR-IOV with MacVTap mode example:

- `interface type='hostdev'`

The `hostdev` type attribute allows you to assign a host device directly to a guest.

- `address type='pci' domain='0' bus='0x09' slot='0x0' function='0x01'`.

The `address type` attribute value specifies the PCI address of the host VF. To obtain the address information, you can run the `lspci` command at the Linux prompt. When you run the command, an address string is displayed, for example, `09:00.1 Ethernet controller: Cisco Systems Inc Device`

0071 (rev a2). In the address string 09.00.1, 09 indicates the bus, 00 indicates the slot, and 1 indicates the function.

- `mac address='01:23:45:67:89:ab'`

Explicit specification of the MAC address is optional. Enter a MAC address that you obtained from your network administrator. If this line is omitted, libvirt generates a MAC address for the vNIC.



Note We recommend that you do not assign a MAC address used by another VM, even if that VM is currently shut down or is no longer used. If you must reuse a MAC address from a previous VM, make sure that the retention timer has expired and ensure that the previous VM is no longer present in the Cisco UCS Manager view.

Example of Using a Network XML File to Specify a Pool of Devices

This example shows how to use a network XML file to specify a pool of devices. In RHEL 6.2 or later, create the network file in `/etc/libvirt/qemu/networks`. List the devices and define a portgroup:

```
<network>
  <name>macvtap_passthru_network</name>
  <forward mode='passthrough'>
    <interface dev='eth2' />
    <interface dev='eth3' />
  </forward>
  <portgroup name='engineering'>
    <virtualport type='802.1Qbh'>
      <parameters profileid='my-port-profile-3' />
    </virtualport>
  </portgroup>
</network>
```

Edit the domain XML file of the VM to reference the network file and portgroup:

```
<domain type='kvm'>
  <name>vm1-rhel6.2</name>
  ...
  <devices>
    ...
    <interface type='network'>
      <mac address='01:23:45:67:89:ab' />
      <source network='macvtap_passthru_network' portgroup='engineering' />
      <model type='virtio' />
    </interface>
    ...
  </devices>
  ...
</domain>
```

After you create the new network XML file, use the `virsh net-define <new-xml-filename>` command to create the new network from the new network XML file.



Tip

You can find the network-related `virsh` commands with `virsh help | grep net-`

You can view help on any `virsh` command with `virsh help <command-name>`

This list describes the components and attributes of the interface element that differ from those described in the SR-IOV with MacVTap mode example:

- `interface type='network'`

The `network` type attribute value specifies an attachment of the VM vNIC to a PCI network device from the pool listed in a separate network file.

- `source network='macvtap_passthru_network' portgroup='engineering'`

The `network` and `portgroup` attribute values specify the name of a network XML file and its pool of network devices.

Activating Intel VT-d in the Kernel

Perform this procedure on the KVM host to enable Intel VT-d extensions, which are required for SR-IOV passthrough.

For more information about this feature in Red Hat Enterprise Linux (RHEL) systems, see the *Virtualization Deployment and Administration Guide*.

Procedure

- Step 1** On the KVM host, open the file `grub.conf` for editing.
The file is typically located in the `/boot` directory. In RHEL systems, you can also access it using the `grub.conf` link in the `/etc` directory.
- Step 2** Locate the line beginning with `kernel`.
- Step 3** Append the command `intel_iommu=on` to the kernel line..

Example:

```
kernel /vmlinuz-2.6.18-190.e15 ro root=/dev/VolGroup00/LogVol100 \  
rhgb quiet intel_iommu=on
```

- Step 4** Save the file.
-

What to Do Next

Reboot the host.



Configuring a Service Profile with VM-FEX

- [Configuring Dynamic vNIC Connection Policies, page 17](#)

Configuring Dynamic vNIC Connection Policies

Dynamic vNIC Connection Policy

**Note**

In an SR-IOV topology, such as a Hyper-V or KVM cluster, a Virtual Function (VF) takes the place of the dynamic vNIC. The VF is essentially a restricted version of the dynamic vNIC, in which all system communication and configuration of the VF is performed through the associated physical function (PF).

The dynamic vNIC connection policy determines how the connectivity between VMs and dynamic vNICs is configured. This policy is required for Cisco UCS domains that include servers with VIC adapters on which you have installed VMs and configured dynamic vNICs.

Ethernet Adapter Policy

Each dynamic vNIC connection policy includes an Ethernet adapter policy and designates the number of vNICs that can be configured for any server associated with a service profile that includes the policy.

For KVM, use the predefined Ethernet adapter policy named Linux.

Server Migration

**Note**

If you migrate a server that is configured with dynamic vNICs, the dynamic interface used by the vNICs fails and notifies you of that failure.

When the server comes back up, assigns new dynamic vNICs to the server. If you are monitoring traffic on the dynamic vNIC, you must reconfigure the monitoring source.

Creating a Dynamic vNIC Connection Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **LAN > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.
- Step 4** Right-click the **Dynamic vNIC Connection Policies** node and choose **Create Dynamic vNIC Connection Policy**.
- Step 5** In the **Create Dynamic vNIC Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the policy.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p> <p>Note Do not specify "default" as the value for the dynamic vNIC connection policy name. Cisco UCS Manager automatically resolves any empty policy references to "default". Any service profiles or service profile templates with only static vNICs defined will automatically reference the policy "default" when it is present. If you specify "default" for the dynamic vNIC connection policy name, then unexpected dynamic vNICs might be created on those service profiles or service profile templates.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
Number of Dynamic vNICs field	<p>The number of dynamic vNICs that this policy affects.</p> <p>Enter an integer between 0 and 256. The default is 54.</p> <p>Note Components of your system might limit this number to fewer than 256 vNICs.</p>
Adapter Policy drop-down list	<p>The adapter profile associated with this policy. The profile must already exist to be included in the drop-down list.</p>

Name	Description
Protection field	Dynamic vNICs are always protected in Cisco UCS, but this field allows you to select a preferred fabric, if any. You can choose one of the following: <ul style="list-style-type: none"> • Protected Pref A—Cisco UCS attempts to use fabric A but fails over to fabric B if necessary • Protected Pref B—Cisco UCS attempts to use fabric B but fails over to fabric A if necessary • Protected—Cisco UCS uses whichever fabric is available

Step 6 Click **OK**.

Step 7 If a confirmation dialog box appears, click **Yes**.

Changing a Dynamic vNIC Connection Policy

Procedure

Step 1 In the **Navigation** pane, click **LAN**.

Step 2 Expand **LAN > Policies**.

Step 3 Expand the node for the organization that contains the policy that you want to change. If the system does not include multitenancy, expand the **root** node.

Step 4 Expand the **Dynamic vNIC Connection Policies** node and click the policy that you want to change.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Change one or more of the following fields:

Name	Description
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Number of Dynamic vNICs field	The number of dynamic vNICs that this policy affects. Enter an integer between 0 and 256. The default is 54. Note Components of your system might limit this number to fewer than 256 vNICs.

Name	Description
Adapter Policy drop-down list	The adapter profile associated with this policy. The profile must already exist to be included in the drop-down list.

You cannot change the other properties of the policy, such as the **Name** field.

- Step 7** Click **Save Changes**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
-

Deleting a Dynamic vNIC Connection Policy

You can delete a dynamic vNIC connection policy.

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **LAN > Policies > Organization_Name**.
- Step 3** Expand the **Dynamic vNIC Connection Policies** node.
- Step 4** Right-click the policy that you want to delete and choose **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-

Viewing Dynamic vNIC Properties in a VM

You can view dynamic vNIC properties in a VM.

Before You Begin

The VM must be operational.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
- Step 2** On the VM tab, expand **All > Clusters**.
- Step 3** Expand **Virtual Machines**.
- Step 4** Expand the virtual machine that contains the dynamic vNIC.
- Step 5** Choose the dynamic vNIC.
- Step 6** In the **Work** pane, click the **General** tab.

In the **Properties** area, the vNIC properties appear.



Configuring Port Profiles

- [Port Profiles](#), page 23
- [Creating a Port Profile](#), page 23
- [Modifying the VLANs in a Port Profile](#), page 25
- [Changing the Native VLAN for a Port Profile](#), page 26
- [Adding a VLAN to a Port Profile](#), page 26
- [Deleting a VLAN from a Port Profile](#), page 26
- [Deleting a Port Profile](#), page 27
- [Port Profile Client](#), page 27

Port Profiles

Port profiles contain the properties and settings that you can use to configure virtual interfaces in Cisco UCS for VM-FEX. The port profiles are created and administered in Cisco UCS Manager. After a port profile is created, assigned to, and actively used by one or more clusters, any changes made to the networking properties of the port profile in Cisco UCS Manager are immediately applied to those clusters.

Creating a Port Profile



Note

In a VM-FEX for KVM system, the following conditions apply:

- The **Max Ports** field applies to the cluster; there is no distributed virtual switch (DVS).
 - The **Host Network IO Performance** field has no effect.
-

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
- Step 2** Expand the **All** node.
- Step 3** Right-click the **Port Profiles** node and choose **Create Port Profile**.
- Step 4** In the **Create Port Profile** dialog box, complete the following fields:

Name	Description
Name field	The user-defined name for the port profile. This name can be between 1 and 31 ASCII alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), and : (colon), and you cannot change this name after the object has been saved.
Description field	The user-defined description for the port profile. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
QoS Policy drop-down list	The quality of service policy associated with this port profile.
Network Control Policy drop-down list	The network control policy associated with this port profile.
Max Ports field	The maximum number of ports that can be associated to a port profile is 4096. The default is 64 ports.
Host Network IO Performance field	This can be one of the following: <ul style="list-style-type: none"> • None—Traffic to and from a virtual machine passes through the DVS. • High Performance— Traffic to and from a virtual machine bypasses the DVS and hypervisor and travels directly between the virtual machines and a virtual interface card (VIC) adapter. <p>In a VM-FEX for KVM system, the Host Network IO Performance field has no effect.</p>
Pin Group drop-down list	The pin group associated with this port profile.
Type field	The type of port profile. This can be one of the following: <ul style="list-style-type: none"> • Regular—Used for all port profiles except for SCVMM. Select regular port profile for KVM and VMware hypervisors. • SLA Only—Used for SCVMM only. <p>Note This field is not displayed for port profile clients used by logical switches.</p>

Step 5 In the **VLANs** area, complete the following fields:

Name	Description
Select column	Check the check box in this column for each VLAN that you want to use.
Name column	The name of the VLAN.
Native VLAN column	To designate one of the VLANs as the native VLAN, click the radio button in this column.

Step 6 Click **OK**.

Modifying the VLANs in a Port Profile

Procedure

Step 1 In the **Navigation** pane, click **VM**.

Step 2 Expand **All > Port Profiles**.

Step 3 Right-click the port profile for which you want to modify the VLANs and choose **Modify VLANs**.

Step 4 In the **Modify VLANs** dialog box, change one or more of the following:

Name	Description
Select column	Check the check box in this column for each VLAN that you want to use. Note VLANs and PVLANS can not be assigned to the same vNIC.
Name column	The name of the VLAN.
Native VLAN column	To designate one of the VLANs as the native VLAN, click the radio button in this column.

Step 5 Click **OK**.

Changing the Native VLAN for a Port Profile

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
 - Step 2** Expand **All > Port Profiles**.
 - Step 3** Right-click the port profile for which you want to change the native VLAN and choose **Modify VLANs**.
 - Step 4** In the **Modify VLANs** dialog box, do the following:
 - a) In the **Native VLAN** column, click the radio button in the row for the VLAN that you want to become the native VLAN.
 - b) Click **OK**.
-

Adding a VLAN to a Port Profile

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
 - Step 2** Expand **All > Port Profiles**.
 - Step 3** Right-click the port profile to which you want to add a VLAN and choose **Modify VLANs**.
 - Step 4** In the **Modify VLANs** dialog box, do the following:
 - a) In the **Select** column, check the check box in the row for the VLAN that you want to add to the port profile.
 - b) (Optional) If you want this VLAN to be the native VLAN, click the radio button in the **Native VLAN** column.
 - c) Click **OK**.
-

Deleting a VLAN from a Port Profile

You can delete a VLAN from a port profile or change the VLAN that you have assigned as the native VLAN through the following steps.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
 - Step 2** Expand **All > Port Profiles**.
 - Step 3** Right-click the port profile from which you want to delete a VLAN and choose **Modify VLANs**.
 - Step 4** In the **Modify VLANs** dialog box, do the following:
 - a) In the **Select** column, uncheck the check box in the row for the VLAN that you want to delete from the port profile.
 - b) (Optional) You can change the native VLAN to a different VLAN by clicking the radio button in the **Native VLAN** column for a different VLAN.
 - c) Click **OK**.
-

Deleting a Port Profile

You cannot delete a port profile if a VM is actively using that port profile.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
 - Step 2** Expand **All > Port Profiles**.
 - Step 3** Right-click the port profile you want to delete and choose **Delete**.
 - Step 4** If a confirmation dialog box displays, click **Yes**.
 - Step 5** Click **OK**.
Cisco UCS Manager deletes the port profile and all its associations.
-

Port Profile Client

Creating a Profile Client

You can create a profile client.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
- Step 2** Expand **All > Port Profiles**.
- Step 3** Right-click the port profile for which you want to create a profile client and choose **Create Profile Client**.
- Step 4** In the **Create Profile Client** dialog box, complete the following fields:

Name	Description
Name field	The user-defined name for the profile client. This name can be between 1 and 16 ASCII alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), and : (colon), and you cannot change this name after the object has been saved.
Description field	The user-defined description of the client. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Distributed Virtual Switch drop-down list	Choose a cluster from the drop-down list.

- Step 5** Click **OK**.

Modifying a Profile Client

You can modify a profile client.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
- Step 2** Expand **All > Port Profiles**.
- Step 3** Click the port profile for which you want to modify the profile client.
- Step 4** In the **Work** pane, click the **Profile Clients** tab.
- Step 5** Right-click the profile client you want to modify and choose **Show Navigator**.
- Step 6** In the Navigator for the profile client, change the values for one or more of the following fields:

Name	Description
Name field	The user-defined name for the profile client.

Name	Description
Description field	The user-defined description of the client. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Datacenter field	A regular expression used to select the appropriate datacenter.
Folder field	A regular expression used to select the appropriate datacenter folder.
Distributed Virtual Switch field	A regular expression used to select the appropriate virtual switch.

Step 7 Click **OK**.

Deleting a Profile Client

You cannot delete a port profile client if a VM is actively using the port profile with which the client is associated.

Procedure

- Step 1** In the **Navigation** pane, click **VM**.
 - Step 2** Expand **All > Port Profiles**.
 - Step 3** Click the port profile from which you want to delete a profile client.
 - Step 4** In the **Work** pane, click the **Profile Clients** tab.
 - Step 5** Right-click the profile client that you want to delete and choose **Delete**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
 - Step 7** Click **Save Changes**.
-

