Introduction

This chapter includes the following sections:

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.

Note All models of the Cisco VICs, including **gen2**, support VM-FEX in the high performance mode as long as you are running ESXi 5.0 or higher on the hypervisor host.

In a VMware environment, VM-FEX supports the standard VMware integration with VMware ESX hypervisors installed on the server and all virtual machine management performed through the VMware vCenter.

Virtualization with Network Interface Cards and Converged Network Adapters

Network interface card (NIC) and converged network adapters support virtualized environments with the standard VMware integration with ESX installed on the server and all virtual machine management performed through the VC.

Portability of Virtual Machines

If you implement service profiles you retain the ability to easily move a server identity from one server to another. After you image the new server, the ESX treats that server as if it were the original.

Communication between Virtual Machines on the Same Server

These adapters implement the standard communications between virtual machines on the same server. If an ESX host includes multiple virtual machines, all communications must go through the virtual switch on the server.

If the system uses the native VMware drivers, the virtual switch is out of the network administrator's domain and is not subject to any network policies. As a result, for example, QoS policies on the network are not applied to any data packets traveling from VM1 to VM2 through the virtual switch.

If the system includes another virtual switch, such as the Nexus 1000, that virtual switch is subject to the network policies configured on that switch by the network administrator.

VM-FEX for VMware Components and Requirements

At a high level, VM-FEX for VMware requires a hypervisor host, Cisco UCS Manager, and VMware vCenter virtualization management software.

The following figure shows these three main components and the methods by which they are connected.

Figure 1: Component Connectivity for VM-FEX with VMware

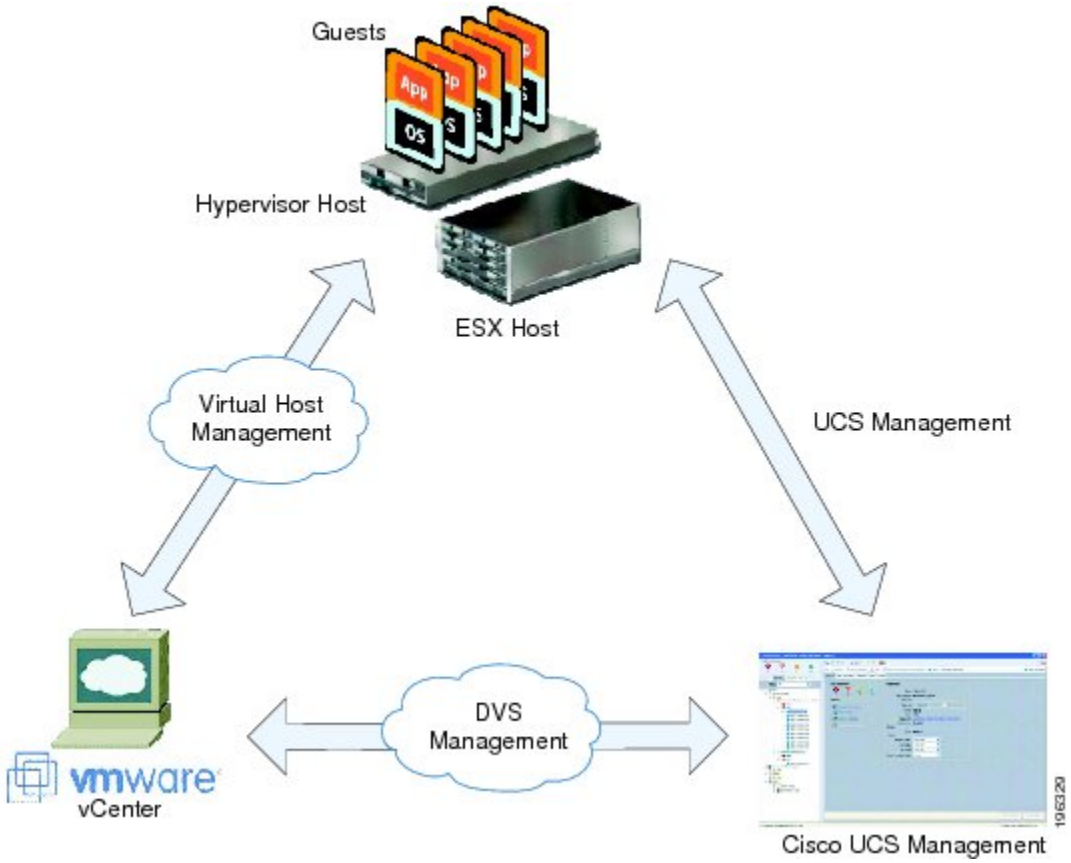

These components must be configured correctly for VM-FEX for VMware to work.

Hypervisor Host

The hypervisor host has these requirements:

- You must install a Cisco UCS Virtual Interface Card in the server you intend to use as the hypervisor host. For more information about installing Cisco UCS Virtual Interface Cards, see the *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.
- You must install the correct version of VMware ESX or ESXi software on the Cisco UCS Manager host. For VM-FEX in standard mode, you must install VMware ESX version 4.0, Update 1 software or later versions.
- You must obtain the Cisco VM-FEX Driver software bundle from the Cisco **Download Software** page:
  - For B-Series from, http://software.cisco.com/download/navigator.html. Select **Downloads Home** > **Products** > **Servers - Unified Computing** > **Cisco UCS B-Series Blade Server Software** and click **Unified Computing Systems (UCS) Drivers**. From this page, click **Latest Releases** or **All Releases** to select the release you want, and click **Download** to download the **ISO image of UCS-related drivers**.
  - For C-Series from, http://software.cisco.com/download/navigator.html. Select **Downloads Home** > **Products** > **Servers - Unified Computing** > **Cisco UCS C-Series Rack-Mount UCS-Managed Server Software** and click **Unified Computing Systems (UCS) Drivers**. From this page, click **Latest Releases** or **All Releases** to select the release you want, and click **Download** to download the **ISO image of UCS-related drivers**.
- You must install the correct version of the Cisco VM-FEX Driver software bundle on the hypervisor host. The Cisco VM-FEX Driver software bundle that you install depends on the VMware ESX or ESXi version you have installed on the hypervisor host. For information about the compatible versions of VMware ESX software and Cisco VM-FEX Driver software bundles, see VM-FEX Software Interoperability Matrix in *Hardware and Software Interoperability Matrix for B Series Servers* at http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html.

Note The VEM software bundle is also a component of another product: the Cisco Nexus 1000V switch. Do not be concerned if you see references to this product during the installation of the VEM bundle. This reference is cosmetic only and does not affect the outcome of the installation and implementation of VM-FEX.

Cisco UCS Manager

VM-FEX for VMware-Related Policies

You must modify or create several policies in order for VM-FEX for VMware to function optimally:

- VMwarePassThrough Ethernet Adapter Policy (high-performance mode only)
- Dynamic vNIC Connection Policies
- BIOS Policy (high-performance mode only)
- VM Lifecycle Policy

Extension File for Communicating with VMware vCenter

For Cisco UCS domains that use VIC adapters to implement VM-FEX, you must create and install an extension file to establish the relationship and communications between Cisco UCS Manager and VMware vCenter. This extension file is an XML file that contains an extension key and public secure sockets layer (SSL) certificate.

Important You cannot change an extension key that is being used by a DVS or an external virtualization manager. If you want to use a custom extension key, we recommend that you create and register the custom key before you create the DVS in Cisco UCS Manager to avoid any possibility of having to delete and recreate the associated DVS.

Extension Key

Cisco UCS and VMware vCenter must be connected for management integration and network communication with the host. To accomplish this connectivity, Cisco UCS provides an extension key that represents the Cisco UCS identity. The extension key must be registered with the external virtualization manager before the Cisco UCS domain can be acknowledged and management and network connectivity can be established.

SSL Certificate

Cisco UCS Manager generates a default, self-signed SSL certificate to support communication with a VMware vCenter. You can also create your own custom certificate to communicate with multiple VMware vCenters. When you create a custom certificate, Cisco UCS Manager recreates the extension files to include the new certificate. If you subsequently delete the custom certificate, Cisco UCS Manager recreates the extension files to include the default, self-signed SSL certificate.

To create a custom certificate, you must obtain and copy an external certificate into Cisco UCS and then create a certificate for VM-FEX that uses the certificate you copied into Cisco UCS.

Distributed Virtual Switches (DVSes)

The Cisco UCS distributed virtual switch (DVS) is a software-based virtual switch that runs along side the vSwitch in the ESX hypervisor and can be distributed across multiple ESX hosts. Unlike the vSwitch, which uses its own local port configuration, a DVS that is associated with multiple ESX hosts uses the same port configuration across all ESX hosts.

After associating an ESX host to a DVS, you can migrate existing VMs from the vSwitch to the DVS, and you can create VMs to use the DVS instead of the vSwitch. With the VM-FEX for VMware implementation, when a VM uses the DVS, all VM traffic passes through the DVS and ASIC-based switching is performed by the fabric interconnect.

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 DVSes, any changes made to the networking properties of the port profile in Cisco UCS Manager are immediately applied to those DVSes.

In VMware vCenter, a port profile is represented as a port group. Cisco UCS Manager pushes the port profile names to VMware vCenter, which displays the names as port groups. None of the specific networking properties or settings in the port profile are visible in VMware vCenter. You must configure at least one port profile client for a port profile if you want Cisco UCS Manager to push the port profile to VMware vCenter.

Port Profile Clients

The port profile client determines the DVSes to which a port profile is applied. By default, the port profile client specifies that the associated port profile applies to all DVSes in the VMware vCenter. However, you can configure the client to apply the port profile to all DVSes in a specific datacenter or datacenter folder or only to one DVS.

VMware vCenter

You need VMware vCenter (vCenter Server and vSphere Client) for VM-FEX for VMware.The VMware vCenter must meet the following requirements:

- The Windows-based machine that you install VMware vCenter on must have network connectivity to the Cisco UCS management port and to the uplink Ethernet port(s) being used by the ESX host. The management port connectivity is used for management plane integration between VMware vCenter and Cisco UCS Manager; the uplink Ethernet port connectivity is used for communication between VMware vCenter and the ESX host.

  Note The HTTP and HTTPS ports (normally TCP 80 and 443) must not be blocked between vCenter and the Cisco UCS domain.

- A VMware vCenter extension key provided by Cisco UCS Manager must be registered with VMware vCenter before VMware vCenter acknowledges the Cisco UCS domain.

Note VMware vCenter Server Appliance (VCSA) is not supported.

In addition, you must configure VMware vCenter with the following parameters:

- A datacenter.
- A distributed virtual switch (DVS).
- ESX hosts added to the DVS and configured to migrate to pass-through switching PTS/DVS.
- For each Cisco VIC adapter, two static vNICs (one for each fabric) added to the DVS.
- Virtual machines (VMs) required for the VMs on the server.
- (For VMware vMotion) Hosts with common shared storage (datastore) that are properly configured for vMotion.
- (For VM-FEX in high-performance mode) All guest memory on the VMs must be reserved.
- (For VM-FEX in high-performance mode) The port profiles and VMwarePassThrough Ethernet adapter policy that you have previously configured in Cisco UCS Manager must be specified.

For information about how to configure these required components in VMware vCenter, see the VMware product documentation.

Modes of Operation

VM-FEX ports can operate in standard mode or high-performance mode.

Standard Mode

In standard mode, traffic to and from a virtual machine passes through the distributed virtual switch (DVS) and the hypervisor.

The following figure shows the traffic paths taken by VM traffic on a Cisco UCS server with a VIC adapter that has VM-FEX ports in standard mode.

Figure 2: Traffic Paths for VM Traffic with VM-FEX

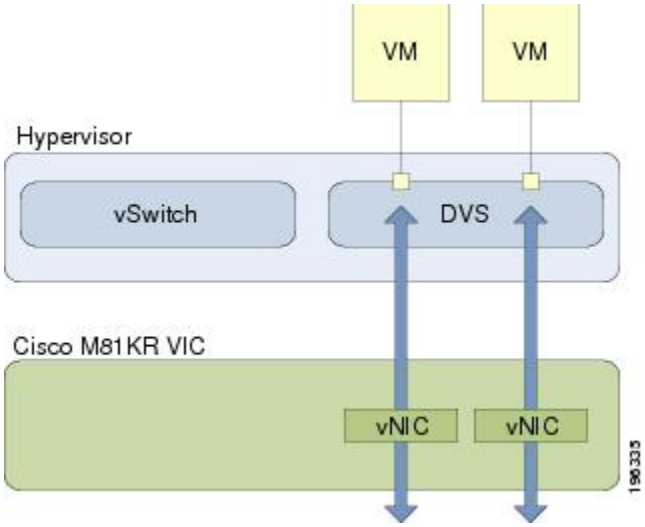

High-Performance Mode

In high-performance mode, traffic to and from a virtual machine (VM) bypasses the DVS and hypervisor. Traffic travels directly between VMs and the virtual interface card (VIC) adapter.

The benefits of high-performance mode are as follows:

- Increases I/O performance and throughput.
- Decreases I/O latency.
- Improves CPU utilization for virualized I/O-intensive applications.

With VMware, high-performance mode also supports vMotion. During vMotion, the hypervisor reconfigures links in high-performance mode to be in standard mode, transitions the link to the new hypervisor, and then

reconfigures the link to be in high-performance mode. The following figure shows how VM-FEX operates in high-performance mode with vMotion.

Figure 3: VM-FEX in High-Performance Mode with VMotion

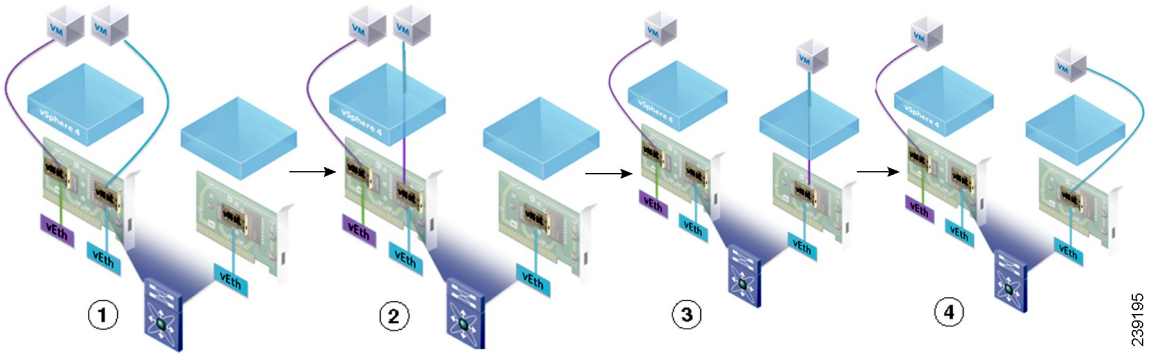

1 Two VMs are attached to a VIC in high-performance mode.

2 VMotion begins on one VM. This VM transitions to standard mode.

3 The VM migrates to the other host, and standard mode is established.

4 The VM transitions back to high-performance mode.

Configuring VM-FEX for VMware

Procedure

	Command or Action	Purpose
Step 1	Configure a Service Profile for VM-FEX for VMware.	You must modify or create several policies in order for VM-FEX for VMware to function optimally: • VMwarePassThrough Ethernet Adapter Policy (high-performance mode only) • Dynamic vNIC Connection Policies • BIOS Policy (high-performance mode only) • VM Lifecycle Policy For more information, see Configuring a Service Profile with VM-FEX.
Step 2	Configure the installation of the Cisco VM-FEX Driver software bundle on the hypervisor host.	You must configure the VMware ESX host and install the Cisco VM-FEX Driver software bundle and a VMware vCenter for VM-FEX. For more information, see Installing the Cisco VM-FEX Driver Software Bundle and the VMware documentation.

	Command or Action	Purpose
Step 3	Connect Cisco UCS Manager to VMware vCenter.	You must connect Cisco UCS Manager with VMware vCenter to manage the distributed virtual switch (DVS) in Cisco UCS Manager. For more information, see Connecting Cisco UCS Manager to VMware vCenter.
Step 4	In Cisco UCS Manager, define a distributed virtual switch.	You must create a distributed virtual switch (DVS) to use in place of the VMware vSwitch. For more information, see Configuring Distributed Virtual Switches in Cisco UCS.
Step 5	In Cisco UCS Manager, define a port profile and (optionally) create a port profile client.	You must create a port profile to define the properties and settings used to configure the virtual interfaces in the DVS. Optionally, you can also create a port profile client that defines the DVSes to which port profiles are assigned. For more information, see Configuring Port Profiles.