



Introduction

- [Overview of Virtualization, on page 1](#)
- [Overview of Cisco Virtual Machine Fabric Extender, on page 1](#)
- [Virtualization with a Virtual Interface Card Adapter, on page 2](#)
- [Single Root I/O Virtualization, on page 2](#)
- [VM-FEX for Hyper-V, on page 3](#)
- [Supported Microsoft Software Versions, on page 5](#)

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 116 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:

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

The Microsoft Hyper-V hypervisor supports SR-IOV.

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

- Cisco UCS Virtual Interface Card 1240
- Cisco UCS Virtual Interface Card 1280
- 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
- Cisco UCS Virtual Interface Card 1387

VM-FEX for Hyper-V

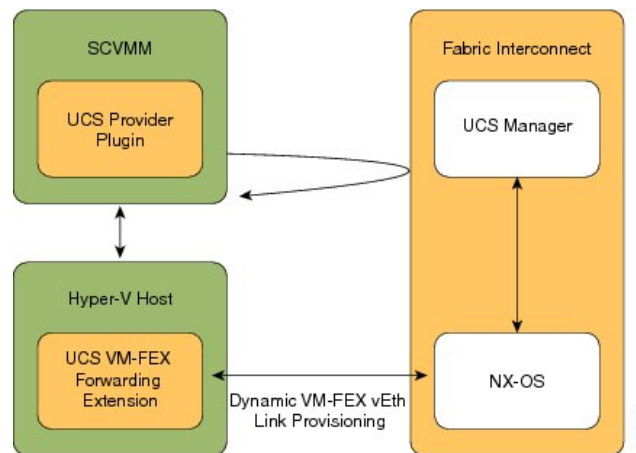
Overview of the Cisco UCS VM-FEX with Microsoft SCVMM Architecture

Cisco UCS Manager (UCSM) and Microsoft System Center Virtual Machine Manager (SCVMM) integration extends the Virtual Machine Fabric Extender (VM-FEX) technology to the Microsoft virtualization platform. The architecture allows Cisco UCSM to configure the networking objects that Microsoft SCVMM uses to set up its networking stacks. Microsoft SCVMM uses the networking objects that are created by Cisco UCSM and deploys them on the Microsoft Hyper-V host that hosts the VMs.

The Hyper-V uses Single Root I/O Virtualization (SR-IOV) technology to deploy virtual connections. The VM's interface gets mapped to the virtual function. The SR-IOV support was added to Cisco UCS Release 2.1 to allow the deployment of VM-FEX in Microsoft Hyper-V hosts, but it lacked a centralized VM network management. Release supports the management plane integration with Microsoft SCVMM and provides a centralized VM network management for the Hyper-V hosts. The deployment leverages the SR-IOV technology that is available on the Cisco virtual interface card (VIC) adapters and enables Cisco UCS fabric interconnects (FIs) to be VM aware.

Figure 1 shows the Cisco UCS VM-FEX with Microsoft SCVMM architecture.

Figure 1: Cisco UCS VM-FEX with Microsoft SCVMM Architecture



Cisco UCSM

Cisco UCSM deploys the service profiles and provisions the baremetal as part of the service profile deployment. While configuring the service profile network settings for the Hyper-V hosts, the administrators have to make

sure that the SR-IOV support is enabled. The network administrator defines the networking objects, for example, the VLANs and the port profiles in Cisco UCSM. These objects get pushed to Cisco NX-OS in the fabric interconnect (FI). The server administrator installs the Cisco UCS provider plugin on Microsoft SCVMM.

Microsoft SCVMM

The Cisco UCS provider plugin enables Microsoft SCVMM to pull the networking objects from Cisco UCSM, use them natively, and deploy them on the Hyper-V hosts. The hosts that are being added to the host groups are the same servers that Cisco UCSM has deployed using the service profiles. It also pulls the network configuration that is specified in Cisco UCSM and pushes it to the Hyper-V host. When you deploy a Logical Switch on the Hyper-V host, the driver extension gets pushed to the host.

The Cisco UCS VM-FEX forwarding extension is a driver extension that is situated on the Hyper-V host. It ensures that the packets are forwarded to the fabric interconnect (FI) and the switching occurs in the FI. The FI is aware of all the MAC addresses of the VMs. The VM-FEX forwarding extension driver gets the configuration from Microsoft SCVMM and instructs Cisco NX-OS to provision a virtual Ethernet interface for virtual NICs (vNICs) that come up on the host.

Dynamic VM-FEX vEth Link Provisioning connects the Hyper-v host and Cisco NX-OS. When a VM is online or when you power on a VM, its network card sends a VIC attach using the Cisco VIC protocol and it gets dynamically connected to the FI.

Hyper-V Host

Microsoft Hyper-V is a virtualization package for Windows Server 2012 and later releases on an x86-64 hardware platform. Hyper-V uses x86 hardware virtualization extensions (for example, Intel VT-x) to implement a hypervisor that hosts VMs as userspace processes.

With VM-FEX for Hyper-V, the 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 of the packets generated by that vNIC.
- For traffic 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 the physical ports, but also between the virtual interfaces (VIFs) that correspond to the vNICs on the VMs.

Networking Terminology

Refer to the following Microsoft networking terminology for more information on the networking objects.

Logical Switch

A logical switch is the native distributed virtual switch (DVS) by Microsoft. It is a template that you can use to instantiate a virtual switch from. You can define a native switch and attach an extension to it. It is known as a switch extension.

Fabric Network

A fabric network is a logical network that has network segments (VLANs) that span across multiple sites. A fabric network can have one or more network sites.

Network Site

A network site includes site-specific network segments. It is also known as a fabric network definition (FND). A network site can have one or more network segments.

Network Segment

A network segment is also known as a VM Network Definition (VMND). It consists of a VLAN and an IP pool.

VM Network

A VM network references a network segment. It is used by the tenant as a network that the network tenants can attach their VMs to. It is the tenant's view of the network.

Virtual Port Profile

A virtual port profile is a profile that defines the quality of service (QoS)/service level agreement (SLA) for a vNIC.

Uplink Port Profile

An uplink port profile carries a list of allowed network segments for a physical network interface card (PNIC).

Reference

For more information on the networking objects, see **Configuring Logical Networking in VMM Overview** at <http://technet.microsoft.com/en-us/library/jj721568.aspx>.

Supported Microsoft Software Versions

Cisco UCS Release supports the following Microsoft software:

- SCVMM 2012 SP1
- Windows Hyper-V 2012 SP1
- SCVMM 2012 R2
- Windows Hyper-V 2012 R2

Reference

For more information on Microsoft SCVMM 2012 SP1 release, see **Description of Update Rollup 4 for System Center 2012 Service Pack 1** at <http://support.microsoft.com/kb/2879276/EN-US>.

For more information on Microsoft SCVMM 2012 R2 release, see **Description of Update Rollup 1 for System Center 2012 R2** at <http://support.microsoft.com/kb/2904734/en-us>.

