

Configuring Single Root I/O Virtualization (SR-IOV)

- Configuring BIOS and SR-IOV VFs, on page 1
- Configuring SR-IOV VFs on the ESXi Host Server , on page 15
- Configuring SR-IOV VFs on the Linux Host Server , on page 21

Configuring BIOS and SR-IOV VFs

Enabling BIOS Parameters

Before you begin

- Ensure your BIOS policy is set up with the following options:
  - For Intel based servers, enable **Intel VT for directed IO** under **Intel Directed IO** tab.

Note Intel VT for directed IO is not available for Intel C220 M8 and Intel C240 M8 platforms.

  - For AMD based servers, enable **IOMMU** and **SVM Mode** under **Processor** tab.

  To update BIOS options, see, Cisco UCS Server BIOS Tokens in Intersight Managed Mode.

- You must have a server profile already created for SR-IOV configuration. To create a Server Profile see Creating a UCS Server Profile. Once the Server Profile is created, follow the steps in this procedure to enable the BIOS policy.

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 Navigate to **Configure** > **Policies** > **Create a Policy**

Step 3 On the **Select Policy Type** page, select **BIOS**, click **Start**.

Step 4 At the **General** page, enter the policy name, and click **Next**.

Step 5 On the **Policy Details** page, configure the following BIOS settings:

- Select **All Platforms**.
- For a server with an Intel CPU, configure the BIOS settings as follows:
  - Enable **Intel VT for Directed IO** under the **Intel Directed IO** drop-down list.
  - Enable **Intel(R) VT** under the **Processor** drop-down list.
- For a server with AMD CPU, configure the BIOS settings as follows:
  - Enable **IOMMU** under **Memory** drop-down list.
  - Enable **SVM Mode** under **Processor** drop-down list.

Step 6 Click **Create**.

Step 7 Associate the BIOS policy to the server profile, and reboot the server.

Note
For more information, see Creating a BIOS Policy and Configuring Server Profile.

Create Ethernet Adapter Policy for SR-IOV

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 In the **Navigation** pane, choose **Configure** > **Policies**, and then click **Create Policy**.

Step 3 Select **Ethernet Adapter**, and then click **Start**.

Step 4 At the **General** page, enter the policy name.

Step 5 Click **Select Cisco Provided Configuration**, choose **SRIOV-HPN**, and click **Select**

Step 6 Click **Next**.

Step 7 Click **Create**.

Enabling SR-IOV VFs using Cisco Intersight GUI

To enable SR-IOV from Cisco Intersight, you must

- Create an SRIOV HPN Connection Policy with desired number of VFs.
- Assign the SRIOV HPN Connection Policy to a Server Profile.

Before you begin

- Ensure that the required BIOS options are enabled before performing this procedure.

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 In the **Navigation** pane, choose **Configure** > **Policies**, and then click **Create Policy**.

Step 3 Select **LAN Connectivity**, and then click **Start**.

Step 4 On the **General** page, enter the following information:

- **Name** of your policy.
- **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.

  A LAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a LAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.
- **Set Tags** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- **Description** to help identify the policy.

Step 5 Click **Next**.

Step 6 On the **Policy Details** page, configure the following:

To set up a vNIC without using a template, click **Add vNIC** and configure the following parameters:

Table 1: For Standalone Servers

Property	Description
Add vNIC Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.	
Name	vNIC name.
Placement Placement Settings for the virtual interface. **Simple** When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.	
Slot ID	When automatic slot ID assignment is disabled, the slot ID needs to be entered manually. Supported values are (1-15) and MLOM.
Uplink Port	Adapter port on which the virtual interface will be created.

Property	Description
PCI link The PCI link used as transport for the virtual interface. **Note** The host device order can get impacted when using both the PCI links and while adding or removing vNICs.	
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs.
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Ethernet Network	Relationship to the Ethernet Network Policy. Select the created Ethernet adapter policy from above. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on Standalone servers. Select or create an Ethernet Network policy.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an Ethernet QoS policy.
Ethernet Adapter	Relationship to the Ethernet Adapter Policy. Select or create an Ethernet Adapter for SR-IOV from above.
Connection	
Disabled	Configuration is disabled.
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.

Property	Description
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. select policy
Class of Service	Class of Service to be used for traffic on the usNIC.
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux 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 vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.

Table 2: For FI-Attached Servers

Property	Description
Enable Azure Stack Host QoS	Enabling AzureStack-Host QoS on an adapter allows the user to carve out traffic classes for RDMA traffic which ensures that a desired portion of the bandwidth is allocated to it.
IQN	
None	This option ensures the IQN name is not associated with the policy.
Pool	
IQN Pool	Relationship to the iSCSI Qualified Name Pool. Select or create an IQN pool.
Static If you select this option, enter a static IQN for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain	
IQN Identifier	User provided static iSCSI Qualified Name (IQN) for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain.
vNIC Configuration	
Manual vNICs Placement	If you select this option, you must manually specify the placement for each vNIC. You can also use the Graphic vNICs Editor to create and specify the placement for each vNIC manually by adding vNICs and slots, and defining the connection between them. **Note** For manual placement, PCI Link is not supported on UCS VIC 1400 Series adapters. If a LAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
Auto vNICs Placement	If you select this option, vNIC placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.
Add vNIC Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.	
Name	Name of the virtual ethernet interface.

Property	Description
Pin Group Name	Pingroup name associated to vNIC for static pinning. LCP deploy will resolve pingroup name and fetches the correspoding uplink port/port channel to pin the vNIC traffic.
MAC	
Pool	If you select this option, select the MAC pool that you want to associate with the LAN Connectivity policy.
MAC Pool	The MAC pool that is assigned. Select or create a MAC pool.
Static	Click **Static** and enter a static MAC address for MAC address assignment. This option is available only for Cisco Intersight Managed FI-Attached servers.
Static MAC Address	The MAC address must be in hexadecimal format xx:xx:xx:xx:xx:xx. To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix 00:25:B5:xx:xx:xx.
Placement	
Simple When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0. **Note** Not applicable for Auto vNIC Placement.	
Switch ID	Refers to the Fabric Interconnect that carries the vNIC traffic.
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The order should start from zero with no overlaps. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. All VIC adapters have a single PCI link except VIC 1340, VIC 1380 and VIC 1385 which have two. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs. Not applicable for Auto vNIC Placement.

Property	Description
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover	Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.
Enabled	Enabling failover ensures that traffic from the vNIC automatically fails over to the secondary Fabric Interconnect, in case the specified Fabric Interconnect path goes down. Failover applies only to Cisco VICs that are connected to a Fabric Interconnect cluster.
Ethernet Network Group	Select or create an Ethernet Network Group policy. You can add multiple Ethernet Network Group Policies (ENGPs) on vNICs. The maximum number of ethernet network group policies is restricted to 50 including shared policies. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. You can associate only one ethernet network group policy with a vNIC if QinQ is configured. The native VLAN must be the same across all ethernet network group policies, or must be set in only one ethernet network group policy. Relationship to the Fabric Ethernet Group Policy. Select or create an ethernet network group policy.
Ethernet Network Control	Relationship to the Fabric Ethernet Network Policy. Select or create an ethernet network control policy.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an ethernet QoS policy.
Ethernet Adapter	Relationship to the the Ethernet Adapter Policy. Select or create an ethernet adapter policy.

Property	Description
iSCSI Boot	Relationship to the boot iSCSI Policy. • Not applicable to SR-IOV. • This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. Select or create an iSCSI boot policy.
Connection	
Disabled	Configuration is disabled.
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. Select or create a usNIC adapter policy.
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux 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 vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	

Property	Description
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.
Placement - Advanced	
Automatic Slot ID Assignment	When enabled, slot ID is determined automatically by the system.
Slot ID	When automatic slot ID assignment is disabled, the slot. ID needs to be entered manually. Supported values are (1-15) and MLOM.
Automatic PCI link Assignment	When enabled, PCI link is determined automatically by the system. **Note** If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vNICs are placed on the same PCI link (link 0). If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vNIC will be placed on PCI link 0.
Load Balanced	When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links. If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to specify the PCI order to load balance the vNICs. If both automatic PCI link assignment and automatic Slot ID are disabled, you need to specify the slot and the PCI order to load balance the vNICs. **Note** You cannot change the PCI link mode of two vNICs from Load Balanced mode to Custom mode without deleting and recreating the vNICs. Enter the following fields for Load Balanced option: Switch ID, PCI Order.

Property	Description
Custom	If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order and the PCI link. **Note** You cannot change the PCI link mode of two vNICs from Custom mode to Load Balanced mode without deleting and recreating the vNICs. Enter the following fields for Custom option: PCI Link, Switch ID, PCI Order.
PCI Link	The PCI Link used as transport for the virtual interface. PCI Link is only applicable for select Cisco UCS VIC 1300 models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored. **Note** Not applicable for Auto vNIC Placement.
Switch ID	The fabric port to which the vNICs will be associated.
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs. Not applicable for Auto vNIC Placement.
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.	

Property	Description
Enabled	Enabling failover ensures that traffic from the vNIC automatically fails over to the secondary Fabric Interconnect, in case the specified Fabric Interconnect path goes down. Failover applies only to Cisco VICs that are connected to a Fabric Interconnect cluster.
Ethernet Network Group	Select or create an Ethernet Network Group policy. You can add multiple Ethernet Network Group Policies (ENGPs) on vNICs. The maximum number of ethernet network group policies is restricted to 50 including shared policies. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. You can associate only one ethernet network group policy with a vNIC if QinQ is configured. The native VLAN must be the same across all ethernet network group policies, or must be set in only one ethernet network group policy. Relationship to the Fabric Ethernet Group Policy. Select or create an ethernet network group policy.
Ethernet Network Control	Relationship to the Fabric Ethernet Network Control Policy. Select or create an ethernet network control policy. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an ethernet QoS policy.
Ethernet Adapter	Relationship to the Ethernet Adapter Policy. Select or create an ethernet adapter policy.
iSCSI Boot	Relationship to the boot iSCSI Policy. • Not applicable to SR-IOV. • This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. Select or create an iSCSI boot policy.
Connection	
Disabled	Configuration is disabled.

Property	Description
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. select policy
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux 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 vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.

Property	Description
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.
Template	Setting up a vNIC using a template.
Add vNIC from Template The source vNIC template to apply to the vNIC instance. All configuration settings from the vNIC template will be applied to the vNIC instance except the overridden list of configurations.	
Name	vNIC name.
vNIC Template	The source vNIC template to apply to the vNIC instance. All configuration settings from the vNIC template will be applied to the vNIC instance except the overridden list of configurations. Select or create an vNIC Template.
Graphics vNIC Editor	Displays the graphics vNIC editor details.

Step 7 Click **Add** and then click **Create**.

Disabling SR-IOV VFs Using Cisco Intersight GUI

Procedure

Step 1 In the **Navigation** pane, click **Policies**.
Step 2 On the **Policies** page, click **Search**.
Step 3 Enter the name of created LAN Connectivity policy from above.
Step 4 Click the **policy**.
Step 5 From **Actions**, select **Edit**.
Step 6 Click **Next**.
Step 7 Select vNIC that you want to disable SR-IOV VFs, and click **Edit**.
Step 8 From **Connection**, click **Disabled**, and then **Update**.
Step 9 Click **Save & Proceed**.

Configuring SR-IOV VFs on the ESXi Host Server

Installing Cisco eNIC Driver

Before you begin

Ensure that the required BIOS parameters and SR-IOV VFs configurations are completed.

The inbox driver does not support SR-IOV functionality. To enable SR-IOV, you must install the appropriate driver. For example, Cisco recommends using the enic drivers for SR-IOV functionality.

Procedure

Step 1 Install the enic driver on the host.

The following example shows the installation of eNIC driver on ESXi:

```
[root@localhost:/vmfs/volumes/645c8bdd-c655e553-8ba0-e8d32272f6c0] esxcli software vib install -v
/vmfs/volumes/C240M7-Standalone/CIS_bootbank_nenic_2.0.15.0-1OEM.800.1.0.20613240.vib --no-sig-check

Installation Result

   Message: The update completed successfully, but the system needs to be rebooted for the changes
to be effective.

   VIBs Installed: CIS_bootbank_nenic_2.0.15.0-1OEM.800.1.0.20613240

   VIBs Removed: CIS_bootbank_nenic_2.0.16.0-1OEM.800.1.0.20613240

   VIBs Skipped:

   Reboot Required: true

   DPU Results:

[root@localhost:/vmfs/volumes/645c8bdd-c655e553-8ba0-e8d32272f6c0]
```

Step 2 Reboot the server to load the enic driver into the running kernel.

Step 3 After reboot, execute the command `esxcli software vib list | grep nenic` to check the driver version.

For more information, see Installing Cisco enic and enic_rdma Drivers.

Verifying the SR-IOV VFs Per Ports on the Host

You can verify the total number of SR-IOV VFs in the following two ways:

Procedure

Step 1 Verify by logging into the VMware ESXi Host Client.:

- Login to the VMware ESXi Host Client.
- Execute the following command to check the vNIC with SR-IOV capability:

```
root@localhost:~] esxcli network sriovnic list
Name    PCI Device    Driver  Link  Speed  Duplex  MAC Address         MTU  Description
-----------------------------------------------------------------------------------------------------
vmnic0  0000:1b:00.0  nenic   Up    50000  Full    f4:ee:31:30:80:40  1500  Cisco Systems Inc
Cisco VIC Ethernet NIC
```

The following output shows the number of VF configured on vNIC:

```
[root@localhost:~] esxcli network sriovnic vf list -n vmnic0
VF ID  Active  PCI Address     Owner World ID
0   false  00000:027:00.1   -
1   false  00000:027:00.2   -
2   false  00000:027:00.3   -
3   false  00000:027:00.4   -
4   false  00000:027:00.5   -
5   false  00000:027:00.6   -
6   false  00000:027:00.7   -
7   false  00000:027:01.0   -
```

Step 2 Alternatively, you can also access your host from vSphere vCenter Client.

For more information on configuring SR-IOV VFs on the host, see Creating SR-IOV VFs on the Host.

After you reboot the host server, do the following:

- Login to the ESXi Host Client, and choose **Networking** > **Virtual Switches**.
- Click **Add Standard Virtual Switch**.
- Add a switch name in the **vSwitch Name** field, select the vmnic with SR-IOV capability, and click **Add**.

  The maximum number of Virtual Functions (VFs) is set to 10.
- In the **Port Groups** tab, click **Add Port Group**.
- In the **Add Port Group** dialog-box, add a new port group and select the switch from the **Virtual Switch** drop-down.

Creating SR-IOV VFs on the Host

Procedure

Step 1 Login to your VMware ESXi Host Client.

Alternatively, you can also access your host from vSphere vCenter Client and browse to **Configure** > **Networking** > **Physical adapters**.

Step 2 Go to **Host** > **Manage** and select the **Hardware** tab.

Step 3 Select **PCI Devices** from the list.

Step 4 From the drop-down list, select **SR-IOV Capable**.

The list shows all the SR-IOV capable devices.

Step 5 Select the vNIC for which you wish to create the VFs.

Step 6 Click **Configure SR-IOV**.

Configure SR-IOV for Cisco VIC Ethernet NIC window is displayed.

Step 7 Perform the following:

Field	Description
Enabled radio button	Select **Yes** to enable the configuration.
Virtual functions field	Number of VFs as configured on SRIOV connection policy that are available for the configuration. Enter an integer between 1 and 64.

Step 8 Click **Save** and then reboot the host server.

Configuring the Switch

Before you begin

Ensure that the SR-IOV VFs are configured.

Procedure

Step 1 Login to your VMware ESXi Host Client.

Step 2 Navigate to **Host** > **Networking** and select the **Virtual switches** tab.

Step 3 Click **Add Standard Virtual Switch**.

Step 4 Enter the name for the switch.

Step 5 Select a SR-IOV Capable Vmnic from the list.

Step 6 Click **Add**.

Step 7 Complete the following:

Field	Description
vSwitch Name field	Enter a suitable name for the virtual switch.
MTU field	Enter the maximum transmission unit. The default is 1500 bytes. The maximum value is 9000 bytes.

Field	Description
Uplink 1 drop-down list	From the drop-down list, select the PCIe devices for which you created the SR-IOVs.
Link Discovery	From the drop-down list, select the **Mode** and the **Protocol**. **Note** These fields remain as default.
Security	Choose from the following options: • **Promiscuous mode**—**Accept**, **Reject**, or **Inherit from vSwitch**. • **MAC address changes**—**Accept**, **Reject**, or **Inherit from vSwitch**. **Forged trasmits**—**Accept**, **Reject**, or **Inherit from vSwitch**.
NIC teaming	Choose from the following: • **Load balancing**—From the drop-down list choose the Load balancing. Values are: **Inherit from vSwitch**, • **Network failover detection**—From the drop-down list choose the network failover detection. Values are: **Inherit from vSwitch**, • **Notify switches**—Choose the notify switches. Values are **Yes**, **No**, **Inherit from vSwitch**. • **Fallback**—Choose the fallback. Values are **Yes**, **No**, **Inherit from vSwitch**. • **Override failover order**—From the drop-down list choose the override failover order. Values are **Yes** or **No**, • **Failover order**—Choose the failover order.
Traffic Shaping	Perform the following: • **Status**—Choose the status. Values are **Enabled**, **Disabled**, **Inherit from vSwitch**. • **Average bandwidth**—Enter the average bandwidth. • **Peek bandwidth**—Enter the peek bandwidth. • **Burst size**—Enter the burst size. **Note** Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the virtual switch.

What to do next

Creating a Virtual Port, on page 19

Creating a Virtual Port

Before you begin

Ensure that the SR-IOV VFs are configured.

Procedure

Step 1 Login to your VMware ESXi Host Client.

Step 2 Go to **Host** > **Networking** and select the **Port Groups** tab.

Step 3 Click **Add port group**.

Add port group-New port group window is displayed

Step 4 Complete the following:

Field	Description
Name field	Enter a suitable name for the virtual port.
VLAN ID field	Enter the VLAN ID.
Virtual Switch drop-down list	From the drop-down list, select the virtual switch.
Security	Choose from the following options: • **Promiscuous mode**—**Accept**, **Reject**, or **Inherit from vSwitch**. • **MAC address changes**—**Accept**, **Reject**, or **Inherit from vSwitch**. **Forged trasmits**—**Accept**, **Reject**, or **Inherit from vSwitch**.

Step 5 Click **Add**.

Creating a New Virtual Machine (VM)

Before you begin

- Login to vCenter using the credentials
- OS ISO image is copied to the datastore of the host server.

Procedure

Installing OS on Guest VM on ESXi.

Adding SR-IOV VF on the Virtual Machine

Before you begin

Power off the Virtual Machine.

Procedure

Step 1 In the Virtual Machine Manager, right-click on the Virtual Machine and select **Open**.

Step 2 Click the **Show Virtual Hardware Detail** icon next to **Monitor** icon.

Step 3 Click **Add Hardware**.

Step 4 In the **Add New Virtual Hardware** window, select **PCI Host Device**. Under the **PCI Device Details** tab, assign a created SR-IOV VF to the Virtual Machine.

Step 5 Click **Finish**.

Step 6 Power on the Virtual Machine.

What to do next

You can now log into the virtual machine, install Cisco eNIC driver version, reboot the virtual machine, and then use the ip link command to verify the added SR-IOV VF. For more information, see Installing Cisco eNIC Driver.

Installing OS on Guest VM on ESXi

Before you begin

Upload the Linux operating system ISO on the datastore.

Procedure

Step 1 Right-click the host node and navigate to **vCenter** > **New Virtual machine**.

Step 2 Select a **Creation Type** > **Create New Virtual Machine**, and click **Next**.

Step 3 Enter a name for the folder, and click **Next**.

Step 4 Select a compute resource, choose a node and click **Next**.

Step 5 Select Storage and check the datastore radio-button, and click **Next**.

Step 6 Select the compatability ESXi 8.0 or later and click **Next**.

Step 7 Select a guest OS version as **RHEL Linux9 (64-bit)**, and click **Next**.

Step 8 Customize the hardware set **CPU** to 2, and **Memory values** to 4 GB.

Step 9 Expand the **Memory** tab, and check **Reserve all guest memory (All locket)** check box.

Step 10 Select **New CD/DVD Drive** *(Datastore ISO file)*, and check the **Connect At Power On** check box.

Step 11 Under **CD/DVD Media**, browse and select the Linux ISO image and click **Next**.

Step 12 Click **Finish**.

Configuring SR-IOV VFs on the Linux Host Server

Installing Cisco eNIC Driver and Enabling IOMMU in Linux Kernel

Before you begin

Ensure that the required BIOS parameters and SR-IOV VFs configurations are completed.

Procedure

Step 1 Install the enic driver on the host.

Following example shows the installation of eNIC driver on RHEL:

```
[user@rack-111 drivers]# rpm -ivh kmod-enic-4.7.0.5-1076.6.rhel9u4_5.14.0_427.13.1.x86_64.rpm
Verifying...                          ################################# [100%]
Preparing...                          ################################# [100%]
Updating / installing...
   1:kmod-enic-4.7.0.5-1076.6.rhel9u4_################################# [100%]
[user@rack-111 drivers]#
```

Step 2 Enable IOMMU on the host using **grubby** command.

Following example shows how to enable IOMMU on RHEL:

```
[user@rack-111 drivers]# grubby --update-kernel=ALL --args="intel_iommu=on iommu=pt"
```

Step 3 Reboot the server to load the enic driver into the running kernel.

Step 4 Execute **modinfo enic** to check enic driver is loaded.

Following example shows the output of **modinfo enic** command:

```
[user@rack-111 drivers]# modinfo enic
filename:       /lib/modules/5.14.0-427.13.1.el9_4.x86_64/extra/enic/enic.ko
version:        4.7.0.5-1076.6
retpoline:      Y
license:        GPL v2
author:         Scott Feldman scofeldm@cisco.com
description:    Cisco VIC Ethernet NIC Driver
rhelversion:    9.4
srcversion:     3A1B1E81C9641925B34D1B2
alias:          pci:v00001137d000002B7sv*sd*bc*sc*i*
alias:          pci:v00001137d00000071sv*sd*bc*sc*i*
```

```
alias:          pci:v00001137d00000044sv*sd*bc*sc*i*
alias:          pci:v00001137d00000043sv*sd*bc*sc*i*
depends:
retpoline:      Y
name:           enic
vermagic:       5.14.0-427.13.1.el9_4.x86_64 SMP preempt mod_unload modversions
sig_id:         PKCS#7
signer:         Cisco UCS Driver Signing REL Cert
sig_key:        D0:54:9A:88:88:DD:0E:7A
sig_hashalgo:   sha256
signature:      89:9C:DA:53:D1:FF:0A:DA:98:9A:7F:AF:63:29:66:EB:FF:0C:D6:65:
                39:6C:15:40:30:6E:99:4B:2C:F0:54:2E:EB:A4:8A:33:D5:9C:41:7A:
                A4:DB:C8:52:55:74:3A:68:F3:22:36:7B:2A:7C:7C:40:8B:7F:6D:9E:
                A5:CF:06:F1:23:42:E6:60:DB:78:0E:46:C9:0C:BC:06:9B:02:A0:AA:
                5A:FC:36:A3:FB:B0:FE:76:F2:EB:2F:AD:AD:84:89:61:30:7D:E9:2F:
                5D:E1:3E:EA:7C:10:B2:42:94:CD:4F:74:19:A6:16:FE:75:B6:78:49:
                E8:F0:4A:A9:01:BB:92:44:A9:FE:C7:CE:DB:E8:F5:08:AF:36:1E:5F:
                30:D3:B1:5F:70:62:56:6F:C2:38:8E:F2:88:28:0F:44:29:E5:44:66:
                34:B7:5C:A7:5E:21:C3:5D:42:D8:C0:87:CA:40:5E:C4:C0:2C:DA:26:
                D2:25:9B:58:A8:84:C6:A6:41:B3:24:9C:D7:E6:4A:79:42:00:32:82:
                7A:CB:36:D8:79:1D:41:1A:9E:1C:A8:0D:39:6D:C8:F1:0D:44:FA:00:
                93:1E:A3:C9:61:AA:DE:25:4A:38:68:C3:9C:14:55:5B:D3:AC:1C:85:
                00:FE:57:F1:DE:F7:A8:04:64:0E:5D:35:D8:AF:CF:A4
parm:           rxcopybreak:Maximum size of packet that is copied to a new buffer on receive (uint)
[user@rack-111 drivers]#
```

Verifying the Total number of SR-IOV VFs per Port on the Host

Before you begin

Ensure that Cisco eNIC driver is installed.

Procedure

Log into the host server and run the following command and replace *interface_name* with actual interface name on the host.

 # cat /sys/class/net/*interface_name*/device/sriov_totalvfs

Example

Following example shows the total number for SR-IOV VFs created from SRIOV HPN Connection Policy on p1p1 interface:

```
[user@rack-111 ~]# cat /sys/class/net/p1p1/device/sriov_totalvfs
32
[user@rack-111 ~]#
```

Creating SR-IOV VFs on the Host

Enabling SR-IOV VFs from SRIOV HPN Connection Policy does not create SR-IOV VFs on the host by default. To create SR-IOV VFs on the host, use the following procedure:

Procedure

Step 1 Execute the following command to create SR-IOV VFs on the host:

 # echo number_of_sriov_devices > /sys/class/net/*sriov interface_name***/device/sriov_numvfs**

Example:

Following example shows the creation of 6 SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# echo 6 > /sys/class/net/p1p1/device/sriov_numvfs
[user@rack-111 ~]#
```

Step 2 Execute the following command to verify the SR-IOV VFs created:

 # cat /sys/class/net/interface_name/device/sriov_numvfs

Example:

Following example shows the verification of SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# cat /sys/class/net/p1p1/device/sriov_numvfs
6
[user@rack-111 ~]#
```

Step 3 (Optional) Alternatively, IP link command shows created SR-IOV VFs.

 # ip link show interface_name

Example:

Following example shows created 6 SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# ip link show pipl
2: p1p1:  <BROADCAST, MULTICAST, UP, LOWER_UP>mtu 9000 qdisc mq state UP mode
DEFAULT group default qlen 10 00
link/ether 98: a2:c0:66:32:80 brd ff:ff:ff:ff:ff:ff
vf 0 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 1 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 2 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 3 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 4 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 5 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
altname enp9s0
altname eno5
[user@rack-111 ~]#
```

Note

After the host server reboots, the created SR-IOV VFs are removed from the host. By adding the command from Step 1 to rc.local file, the same number of SR-IOV VFs can be created each time the host server boots up.

What to do next

You can create a new virtual machine.

Creating a New Virtual Machine (VM)

Before you begin

- Host with Desktop Environment
- Virtualization packages are installed
- Copy OS ISO image to the sever datastore

Procedure

Step 1 Verify the virtualization is enabled on the host server by using this command.

 # lscpu | grep Virtualization

Example:

This example shows the Intel's virtualization technology VT-x is enabled.

```
[user@rack-111 ~]$ lscpu | grep Virtualization
Virtualization: VT-x
[user@rack-111 ~]$
```

Step 2 Verify the KVM modules are loaded by using this command.

 # lsmod | grep kvm

Example:

This example shows KVM modules are loaded in the host server.

```
[user@rack-111 ~]$ lsmod | grep kvm
kvm_intel    409600    8
kvm         1134592    1 kvm_intel
irqbypass      6384    290 vfio_pci_core, kvm
[user@rack-111 ~]$
```

Step 3 Type **virt-manager** command at the terminal to launch Virtual Machine Manager GUI.

Step 4 At the Virtual Machine Manager, click **File** > **New Virtual Machine** to create a new virtual machine.

Step 5 At **New VM window**, select **Local install media (ISO image or CDROM)** option and click **Forward**.

Step 6 At **Choose ISO or CDROM install media**, click **Browse**.

Step 7 At **Locate ISO media volume** window, click **Browser Local**.

Step 8 Go to the folder that has ISO image. Select ISO image and click **Open**.

Step 9 Click **Forward**.

Step 10 Select the desire Memory and CPU settings for the VM and click **Forward**.

Step 11 Choose the VM's disk image size and click **Forward**.

Step 12 Enter a name for the VM in the **Name** field and click **Finish**.

You may monitor the OS installation progress.

Adding SR-IOV VF on the Virtual Machine

Before you begin

Power off the Virtual Machine.

Procedure

Step 1 In the Virtual Machine Manager, right-click on the Virtual Machine and select **Open**.

Step 2 Click the **Show Virtual Hardware Detail** icon next to **Monitor** icon.

Step 3 Click **Add Hardware**.

Step 4 In the **Add New Virtual Hardware** window, select **PCI Host Device**. Under the **PCI Device Details** tab, assign a created SR-IOV VF to the Virtual Machine.

Step 5 Click **Finish**.

Step 6 Power on the Virtual Machine.

What to do next

You can now log into the virtual machine, install Cisco eNIC driver version, reboot the virtual machine, and then use the ip link command to verify the added SR-IOV VF. For more information, see Installing Cisco eNIC Driver.