



Configuring Network-Related Policies

This chapter includes the following sections:

- [Configuring vNIC Templates, page 1](#)
- [Configuring Ethernet Adapter Policies, page 8](#)
- [Configuring the Default vNIC Behavior Policy, page 19](#)
- [Configuring LAN Connectivity Policies, page 20](#)
- [Configuring Network Control Policies, page 26](#)
- [Configuring Multicast Policies, page 29](#)
- [Configuring LACP policies, page 31](#)
- [Configuring UDLD Link Policies, page 32](#)
- [Configuring VMQ Connection Policies, page 37](#)
- [NetQueue, page 42](#)

Configuring vNIC Templates

vNIC Template

The vNIC LAN connectivity policy defines how a vNIC on a server connects to the LAN.

Cisco UCS Manager does not automatically create a VM-FEX port profile with the correct settings when you create a vNIC template. If you want to create a VM-FEX port profile, you must configure the target of the vNIC template as a VM. You must include this policy in a service profile for it to take effect.

You can select VLAN groups in addition to any individual VLAN while creating a vNIC template.

**Note**

If your server has two Emulex or QLogic NICs (Cisco UCS CNA M71KR-E or Cisco UCS CNA M71KR-Q), you must configure vNIC policies for both adapters in your service profile to get a user-defined MAC address for both NICs. If you do not configure policies for both NICs, Windows still detects both of them in the PCI bus. Then because the second eth is not part of your service profile, Windows assigns it a hardware MAC address. If you then move the service profile to a different server, Windows sees additional NICs because one NIC did not have a user-defined MAC address.

Creating a vNIC Template

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- Named VLAN
- MAC pool
- QoS policy
- LAN pin group
- Statistics threshold 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 **vNIC Templates** node and choose **Create vNIC Template**.

Step 5 In the **Create vNIC Template** dialog box:

a) In the **General** area, complete the following fields:

Name	Description
Name field	The name of the vNIC template. 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.
Description field	A user-defined description of the template. 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).

Name	Description
Fabric ID field	<p>The fabric interconnect associated with the component.</p> <p>If you want vNICs created from this template to be able to access the second fabric interconnect if the default one is unavailable, check the Enable Failover check box.</p> <p>Note Do not enable vNIC fabric failover under the following circumstances:</p> <ul style="list-style-type: none"> • If the Cisco UCS domain is running in Ethernet switch mode. vNIC fabric failover is not supported in that mode. If all Ethernet uplinks on one fabric interconnect fail, the vNICs do not fail over to the other. • If you plan to associate one or more vNICs created from this template to a server with an adapter that does not support fabric failover, such as the Cisco UCS 82598KR-CI 10-Gigabit Ethernet Adapter. If so, Cisco UCS Manager generates a configuration fault when you associate the service profile with the server.
Redundancy Type	<p>The Redundancy type that you choose initiates a fabric failover using vNIC/HBA redundancy pairs.</p> <ul style="list-style-type: none"> • Primary Template— Creates configurations that can be shared with the Secondary template. Any other shared changes on the Primary template are automatically synchronized to the Secondary template. • Secondary Template— All shared configurations are inherited from the Primary template. • No Redundancy— Legacy vNIC/vNHBA template behavior. Select this option if you do not want to use redundancy.
Target list box	<p>A list of the possible targets for vNICs created from this template. The target you choose determines whether or not Cisco UCS Manager automatically creates a VM-FEX port profile with the appropriate settings for the vNIC template. This can be one of the following:</p> <ul style="list-style-type: none"> • Adapter—The vNICs apply to all adapters. No VM-FEX port profile is created if you choose this option. • VM—The vNICs apply to all virtual machines. A VM-FEX port profile is created if you choose this option.

Name	Description
Template Type field	<ul style="list-style-type: none"> • Initial Template: vNICs created from this template are not updated if the template changes. • Updating Template: vNICs created from this template are updated if the template changes.

- b) In the **VLANs** area, use the table to select the VLAN to assign to vNICs created from this template. The table contains the following columns:

Name	Description
Select column	<p>Check the check box in this column for each VLAN that you want to use.</p> <p>Note VLANs and PVLANS can not be assigned to the same vNIC.</p>
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.

- c) In the **Policies** area, complete the following fields:

Name	Description
CDN Source field	<p>This can be one of the following options:</p> <ul style="list-style-type: none"> • vNIC Name —Uses the vNIC template name of the vNIC instance as the CDN name. This is the default option. • User Defined — Displays the CDN Name field for you to enter a user-defined CDN name for the vNIC template.
MTU field	<p>The maximum transmission unit, or packet size, that vNICs created from this vNIC template should use.</p> <p>Enter an integer between 1500 and 9000.</p> <p>Note If the vNIC template has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets may be dropped during data transmission.</p>
MAC Pool drop-down list	The MAC address pool that vNICs created from this vNIC template should use.

Name	Description
QoS Policy drop-down list	The quality of service policy that vNICs created from this vNIC template should use.
Network Control Policy drop-down list	The network control policy that vNICs created from this vNIC template should use.
Pin Group drop-down list	The LAN pin group that vNICs created from this vNIC template should use.
Stats Threshold Policy drop-down list	The statistics collection policy that vNICs created from this vNIC template should use.

Step 6 Click **OK**.

What to Do Next

Include the vNIC template in a service profile.

Creating vNIC Template Pairs

Procedure

- Step 1** In the Navigation pane, click the **LAN** tab. On the **LAN** tab, expand **LAN > Policies**.
- Step 2** Expand the node for the organization where you want to create the policy. If the system does not include multi-tenancy, expand the root node.
- Step 3** Right-click the **vNIC Templates** node and choose **Create vNIC Template**. In the **Create vNIC Template** dialog box, assign a **Name**, **Description**, and select the **Fabric ID** for the template.
- Step 4** Select the **Redundancy Type** as **Primary** or **Secondary** or **No Redundancy**. See the redundancy type descriptions below.
- Step 5** Select the **Peer Redundancy Template**—to choose the name of the corresponding **Primary** or **Secondary** redundancy template to perform the template pairing from the **Primary** or **Secondary** redundancy template.
 - **Primary**—Creates configurations that can be shared with the Secondary template. Any other shared changes on the Primary template are automatically synchronized to the Secondary template.
 - **VLANs**
 - **Template Type**
 - **MTU**
 - **Network Control Policies**
 - **Connection Policies**

- QoS Policy
- Stats Threshold Policy

Note

Following is a list of non-shared configurations:

- **Fabric ID**

Note The Fabric ID must be mutually exclusive. If you assign the Primary template to Fabric A, then Fabric B is automatically assigned to the Secondary template as part of the synchronization from the Primary template.

- CDN Source
- MAC Pool
- Description
- Pin Group Policy

- **Secondary—**

All shared configurations are inherited from the Primary template.

- **No Redundancy—**

Legacy vNIC template behavior.

Step 6 Click **OK**.**What to Do Next**

After you create the vNIC redundancy template pair, you can use the redundancy template pair to create redundancy vNIC pairs for any service profile in the same organization or sub-organization.

Undo vNIC Template Pairs

You can undo the vNIC template pair by changing the Peer Redundancy Template so that there is no peer template for the Primary or the Secondary template. When you undo a vNIC template pair, the corresponding vNIC pairs also becomes undone.

Procedure

Select **not set** from the **Peer Redundancy Template** drop-down list to undo the pairing between the peer Primary or Secondary redundancy template used to perform the template pairing. You can also select **None** as the **Redundancy Type** to undo the pairing.

Note If you delete one template in a pair, you are prompt to delete the other template in the pair. If you do not delete the other template in the pair, that template resets its peer reference and retains its redundancy type.

Binding a vNIC to a vNIC Template

You can bind a vNIC associated with a service profile to a vNIC template. When you bind the vNIC to a vNIC template, Cisco UCS Manager configures the vNIC with the values defined in the vNIC template. If the existing vNIC configuration does not match the vNIC template, Cisco UCS Manager reconfigures the vNIC. You can only change the configuration of a bound vNIC through the associated vNIC template. You cannot bind a vNIC to a vNIC template if the service profile that includes the vNIC is already bound to a service profile template.

**Important**

If the vNIC is reconfigured when you bind it to a template, Cisco UCS Manager reboots the server associated with the service profile.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile with the vNIC you want to bind. If the system does not include multi-tenancy, expand the **root** node.
- Step 4** Expand ***Service_Profile_Name*** > **vNICs**.
- Step 5** Click the vNIC you want to bind to a template.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Bind to a Template**.
- Step 8** In the **Bind to a vNIC Template** dialog box, do the following:
 - a) From the **vNIC Template** drop-down list, choose the template to which you want to bind the vNIC.
 - b) Click **OK**.
- Step 9** In the warning dialog box, click **Yes** to acknowledge that Cisco UCS Manager may need to reboot the server if the binding causes the vNIC to be reconfigured.

Unbinding a vNIC from a vNIC Template

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile with the vNIC you want to unbind. If the system does not include multi-tenancy, expand the **root** node.

- Step 4** Expand *Service_Profile_Name* > vNICs.
- Step 5** Click the vNIC you want to unbind from a template.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Actions** area, click **Unbind from a Template**.
- Step 8** If a confirmation dialog box displays, click **Yes**.
-

Deleting a vNIC Template

Procedure

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

Configuring Ethernet Adapter Policies

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter, including how the adapter handles traffic. For example, you can use these policies to change default settings for the following:

- Queues
- Interrupt handling
- Performance enhancement
- RSS hash
- Failover in a cluster configuration with two fabric interconnects

**Note**

For Fibre Channel adapter policies, the values displayed by Cisco UCS Manager may not match those displayed by applications such as QLogic SANsurfer. For example, the following values may result in an apparent mismatch between SANsurfer and Cisco UCS Manager:

- **Max LUNs Per Target**—SANsurfer has a maximum of 256 LUNs and does not display more than that number. Cisco UCS Manager supports a higher maximum number of LUNs.
- **Link Down Timeout**—In SANsurfer, you configure the timeout threshold for link down in seconds. In Cisco UCS Manager, you configure this value in milliseconds. Therefore, a value of 5500 ms in Cisco UCS Manager displays as 5s in SANsurfer.
- **Max Data Field Size**—SANsurfer has allowed values of 512, 1024, and 2048. Cisco UCS Manager allows you to set values of any size. Therefore, a value of 900 in Cisco UCS Manager displays as 512 in SANsurfer.
- **LUN Queue Depth**—The LUN queue depth setting is available for Windows system FC adapter policies. Queue depth is the number of commands that the HBA can send and receive in a single transmission per LUN. Windows Storport driver sets this to a default value of 20 for physical miniports and to 250 for virtual miniports. This setting adjusts the initial queue depth for all LUNs on the adapter. Valid range for this value is 1 to 254. The default LUN queue depth is 20. This feature only works with Cisco UCS Manager version 3.1(2) and higher.
- **IO TimeOut Retry**—When the target device is not responding to an IO request within the specified timeout, the FC adapter will abort the pending command then resend the same IO after the timer expires. The FC adapter valid range for this value is 1 to 59 seconds. The default IO retry timeout is 5 seconds. This feature only works with Cisco UCS Manager version 3.1(2) and higher.

Operating System Specific Adapter Policies

By default, Cisco UCS provides a set of Ethernet adapter policies and Fibre Channel adapter policies. These policies include the recommended settings for each supported server operating system. Operating systems are sensitive to the settings in these policies. Storage vendors typically require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.

**Important**

We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

Completion Queues = Transmit Queues + Receive Queues

Interrupt Count = (Completion Queues + 2) rounded up to nearest power of 2

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

Completion Queues = 1 + 8 = 9

Interrupt Count = (9 + 2) rounded up to the nearest power of 2 = 16

Accelerated Receive Flow Steering

Accelerated Receive Flow Steering (ARFS) is hardware-assisted receive flow steering that can increase CPU data cache hit rate by steering kernel level processing of packets to the CPU where the application thread consuming the packet is running.

Using ARFS can improve CPU efficiency and reduce traffic latency. Each receive queue of a CPU has an interrupt associated with it. You can configure the Interrupt Service Routine (ISR) to run on a CPU. The ISR moves the packet from the receive queue to the backlog of one of the current CPUs, which processes the packet later. If the application is not running on this CPU, the CPU must copy the packet to non-local memory, which adds to latency. ARFS can reduce this latency by moving that particular stream to the receive queue of the CPU on which the application is running.

ARFS is disabled by default and can be enabled through Cisco UCS Manager. To configure ARFS, do the following:

- 1 Create an adapter policy with ARFS enabled.
- 2 Associate the adapter policy with a service profile.
- 3 Enable ARFS on a host.
 - 1 Turn off Interrupt Request Queue (IRQ) balance.
 - 2 Associate IRQ with different CPUs.
 - 3 Enable ntuple by using ethtool.

Guidelines and Limitations for Accelerated Receive Flow Steering

- ARFS supports 64 filters per vNIC
- ARFS is supported on the following adapters:
 - Cisco UCS VIC 1280, 1240, 1340, and 1380
 - Cisco UCS VIC 1225, 1225T, 1285, 1223, 1227, 1227T, 1385, 1387
- ARFS is supported on the following Operating Systems:
 - Red Hat Enterprise Linux 6.5, and 6.6
 - Red Hat Enterprise Linux 7.0 and higher versions
 - SUSE Linux Enterprise Server 11 SP2 and SP3
 - SUSE Linux Enterprise Server 12 and higher versions
 - Ubuntu 14.04.2

Interrupt Coalescing

Adapters typically generate a large number of interrupts that a host CPU must service. Interrupt coalescing reduces the number of interrupts serviced by the host CPU. This is done by interrupting the host only once for multiple occurrences of the same event over a configurable coalescing interval.

When interrupt coalescing is enabled for receive operations, the adapter continues to receive packets, but the host CPU does not immediately receive an interrupt for each packet. A coalescing timer starts when the first packet is received by the adapter. When the configured coalescing interval times out, the adapter generates one interrupt with the packets received during that interval. The NIC driver on the host then services the multiple packets that are received. Reduction in the number of interrupts generated reduces the time spent by the host CPU on context switches. This means that the CPU has more time to process packets, which results in better throughput and latency.

Adaptive Interrupt Coalescing

Due to the coalescing interval, the handling of received packets adds to latency. For small packets with a low packet rate, this latency increases. To avoid this increase in latency, the driver can adapt to the pattern of traffic flowing through it and adjust the interrupt coalescing interval for a better response from the server.

Adaptive interrupt coalescing (AIC) is most effective in connection-oriented low link utilization scenarios including email server, databases server, and LDAP server. It is not suited for line-rate traffic.

Guidelines and Limitations for Adaptive Interrupt Coalescing

- Adaptive Interrupt Coalescing (AIC) does not provide any reduction in latency when the link utilization is more than 80 percent.
- Enabling AIC disables static coalescing.
- AIC is supported on the following Operating Systems:
 - Red Hat Enterprise Linux 6.4 and higher versions
 - Red Hat Enterprise Linux 7.0 and higher versions
 - SUSE Linux Enterprise Server 11 SP2 and SP3
 - SUSE Linux Enterprise Server 12
 - XenServer 6.5
 - Ubuntu 14.04.2

RDMA Over Converged Ethernet for SMB Direct

RDMA over Converged Ethernet (RoCE) allows direct memory access over an Ethernet network. RoCE is a link layer protocol, and hence, it allows communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization and higher utilization of network bandwidth. Windows 2012 R2 and later versions use RDMA for accelerating and improving the performance of SMB file sharing and Live Migration.

Cisco UCS Manager Release 2.2(4) supports RoCE for Microsoft SMB Direct. It sends additional configuration information to the adapter while creating or modifying an Ethernet adapter policy.

Guidelines and Limitations for SMB Direct with RoCE

- Microsoft SMB Direct with RoCE is supported:

- on Windows 2012 R2 for Cisco UCS Manager release 2.2(4) and later releases.
- on Windows 2016 for Cisco UCS Manager release 2.2(8) and later releases.
- Microsoft SMB Direct with RoCE is supported only with third generation Cisco UCS VIC 1340, 1380, 1385, 1387 adapters. Second generation UCS VIC 1225 and 1227 adapters are not supported.
- RoCE configuration is supported between Cisco adapters. Interoperability between Cisco adapters and third party adapters is not supported.
- Cisco UCS Manager does not support more than 4 RoCE-enabled vNICs per adapter.
- Cisco UCS Manager does not support RoCE with NVGRE, VXLAN, NetFlow, VMQ, or vNIC.
- Maximum number of queue pairs per adapter is 8192.
- Maximum number of memory regions per adapter is 524288.
- If you do not disable RoCE before downgrading Cisco UCS Manager from Release 2.2(4), downgrade will fail.
- Cisco UCS Manager does not support fabric failover for vNICs with RoCE enabled.

Creating an Ethernet Adapter Policy



Tip

If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > 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 **Adapter Policies** and choose **Create Ethernet Adapter Policy**.
- Step 5** Enter a **Name** and optional **Description** for the policy.
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.
- Step 6** (Optional) In the **Resources** area, adjust the following values:

Name	Description
Transmit Queues field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.

Name	Description
Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.
Receive Queues field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.
Completion Queues field	The number of completion queue resources to allocate. In general, the number of completion queue resources you should allocate is equal to the number of transmit queue resources plus the number of receive queue resources. Enter an integer between 1 and 512.
Interrupts field	The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources. Enter an integer between 1 and 514.

Step 7 (Optional) In the **Options** area, adjust the following values:

Name	Description
Transmit Checksum Offload field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU calculates all packet checksums. • Enabled—The CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. <p>Note This option affects only packets sent from the interface.</p>
Receive Checksum Offload field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU validates all packet checksums. • Enabled—The CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. <p>Note This option affects only packets received by the interface.</p>

Name	Description
TCP Segmentation Offload field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU segments large TCP packets. • Enabled—The CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. <p>Note This option is also known as Large Send Offload (LSO) and affects only packets sent from the interface.</p>
TCP Large Receive Offload field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU processes all large packets. • Enabled—The hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput. <p>Note This option affects only packets received by the interface.</p>
Receive Side Scaling field	<p>RSS distributes network receive processing across multiple CPUs in multiprocessor systems. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Network receive processing is always handled by a single processor even if additional processors are available. • Enabled—Network receive processing is shared across processors whenever possible.
Accelerated Receive Flow Steering field	<p>Packet processing for a flow must be performed on the local CPU. This is supported for Linux operating systems only. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU is not specified. • Enabled—Packet processing is performed on the local CPU.
Network Virtualization using Generic Routing Encapsulation field	<p>Whether NVGRE overlay hardware offloads for TSO and checksum are enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—NVGRE overlay hardware offloads are not enabled. • Enabled—NVGRE overlay hardware offloads are enabled.
Virtual Extensible LAN field	<p>Whether VXLAN overlay hardware offloads for TSO and checksum are enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—VXLAN overlay hardware offloads are not enabled. • Enabled—VXLAN overlay hardware offloads are enabled.

Name	Description
Failback Timeout field	After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC. Enter a number of seconds between 0 and 600.
Interrupt Mode field	The preferred driver interrupt mode. This can be one of the following: <ul style="list-style-type: none"> • MSI X—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—MSI only. • INTx—PCI INTx interrupts.
Interrupt Coalescing Type field	This can be one of the following: <ul style="list-style-type: none"> • Min—The system waits for the time specified in the Interrupt Timer field before sending another interrupt event. • Idle—The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the Interrupt Timer field.
Interrupt Timer field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter a value between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
RoCE field	Whether Remote Direct Memory Access over an Ethernet network is enabled. This can be one of the following: <ul style="list-style-type: none"> • Disabled—RoCE is disabled on the Ethernet adapter. • Enabled—RoCE is enabled on the Ethernet adapter.
RoCE Properties area	Lists the RoCE properties. This area is enabled only if you enable RoCE.
Queue Pairs	The number of queue pairs per adapter. Enter an integer between 1 and 8192. It is recommended that this number be an integer power of 2.
Memory Regions	The number of memory regions per adapter. Enter an integer between 1 and 524288. It is recommended that this number be an integer power of 2.

Name	Description
Resource Groups	<p>The number of resource groups per adapter.</p> <p>Enter an integer between 1 and 128.</p> <p>It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.</p>

Step 8 Click **OK**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Configuring an Ethernet Adapter Policy to Enable eNIC Support for MRQS on Linux Operating Systems

Cisco UCS Manager includes eNIC support for the Multiple Receive Queue Support (MRQS) feature on Red Hat Enterprise Linux Version 6.x and SUSE Linux Enterprise Server Version 11.x.

Procedure

Step 1 Create an Ethernet adapter policy.
Use the following parameters when creating the Ethernet adapter policy:

- Transmit Queues = 1
- Receive Queues = n (up to 8)
- Completion Queues = # of Transmit Queues + # of Receive Queues
- Interrupts = # Completion Queues + 2
- Receive Side Scaling (RSS) = Enabled
- Interrupt Mode = Msi-X

See [Creating an Ethernet Adapter Policy](#), on page 12.

Step 2 Install an eNIC driver Version 2.1.1.35 or later.
See [Cisco UCS Virtual Interface Card Drivers for Linux Installation Guide](#).

Step 3 Reboot the server

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE

Cisco UCS Manager supports stateless offloads with NVGRE only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running Windows Server 2012 R2 operating systems. Stateless offloads with NVGRE cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > 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 **Adapter Policies** and choose **Create Ethernet Adapter Policy**.
- a) In the **Resources** area, set the following options:
 - Transmit Queues = 1
 - Receive Queues = n (up to 8)
 - Completion Queues = # of Transmit Queues + # of Receive Queues
 - Interrupts = # Completion Queues + 2
 - b) In the **Options** area, set the following options:
 - Network Virtualization using Generic Routing Encapsulation = Enabled
 - Interrupt Mode = Msi-X

For more information on creating an Ethernet adapter policy, see [Creating an Ethernet Adapter Policy](#), on page 12.

- Step 5** Click **OK** to create the Ethernet adapter policy.
- Step 6** Install an eNIC driver Version 3.0.0.8 or later.
For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows/b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.html.
- Step 7** Reboot the server.
-

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with VXLAN

Cisco UCS Manager supports stateless offloads with VXLAN only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running VMWare ESXi Release 5.5 and later releases of the operating system. Stateless offloads with VXLAN cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

-
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > 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 **Adapter Policies** and choose **Create Ethernet Adapter Policy**.
- a) In the **Resources** area, set the following options:
 - Transmit Queues = 1
 - Receive Queues = n (up to 8)
 - Completion Queues = # of Transmit Queues + # of Receive Queues
 - Interrupts = # Completion Queues + 2
 - b) In the **Options** area, set the following options:
 - Virtual Extensible LAN = Enabled
 - Interrupt Mode = Msi-X

For more information on creating an ethernet adapter policy, see [Creating an Ethernet Adapter Policy](#), on page 12.

- Step 5** Click **OK** to create the Ethernet adapter policy.
- Step 6** Install an eNIC driver Version 2.1.2.59 or later.
For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/ESX/2-0/b_Cisco_VIC_Drivers_for_ESX_Installation_Guide.html.
- Step 7** Reboot the server.
-

Deleting an Ethernet Adapter Policy

Procedure

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

Configuring the Default vNIC Behavior Policy

Default vNIC Behavior Policy

Default vNIC behavior policy allows you to configure how vNICs are created for a service profile. You can choose to create vNICs manually, or you can create them automatically.

You can configure the default vNIC behavior policy to define how vNICs are created. This can be one of the following:

- **None**—Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created.
- **HW Inherit**—If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter installed in the server associated with the service profile.

**Note**

If you do not specify a default behavior policy for vNICs, **HW Inherit** is used by default.

Configuring a Default vNIC Behavior Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **LAN > Policies**.
- Step 3** Expand the **root** node.
You can configure only the default vNIC behavior policy in the root organization. You cannot configure the default vNIC behavior policy in a sub-organization.
- Step 4** Click **Default vNIC Behavior**.
- Step 5** On the **General Tab**, in the **Properties** area, click one of the following radio buttons in the **Action** field:
 - **None**—Cisco UCS Manager does not create default vNICs for a service profile. All vNICs must be explicitly created.
 - **HW Inherit**—If a service profile requires vNICs and none have been explicitly defined, Cisco UCS Manager creates the required vNICs based on the adapter installed in the server associated with the service profile.
- Step 6** Click **Save Changes**.

Configuring LAN Connectivity Policies

About the LAN and SAN Connectivity Policies

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNS to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.

**Note**

We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates and can be used to configure multiple servers.

Privileges Required for LAN and SAN Connectivity Policies

Connectivity policies enable users without network or storage privileges to create and modify service profiles and service profile templates with network and storage connections. However, users must have the appropriate network and storage privileges to create connectivity policies.

Privileges Required to Create Connectivity Policies

Connectivity policies require the same privileges as other network and storage configurations. For example, you must have at least one of the following privileges to create connectivity policies:

- admin—Can create LAN and SAN connectivity policies
- ls-server—Can create LAN and SAN connectivity policies
- ls-network—Can create LAN connectivity policies
- ls-storage—Can create SAN connectivity policies

Privileges Required to Add Connectivity Policies to Service Profiles

After the connectivity policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create connectivity policies.

Interactions between Service Profiles and Connectivity Policies

You can configure the LAN and SAN connectivity for a service profile through either of the following methods:

- LAN and SAN connectivity policies that are referenced in the service profile
- Local vNICs and vHBAs that are created in the service profile
- Local vNICs and a SAN connectivity policy
- Local vHBAs and a LAN connectivity policy

Cisco UCS maintains mutual exclusivity between connectivity policies and local vNIC and vHBA configuration in the service profile. You cannot have a combination of connectivity policies and locally created vNICs or vHBAs. When you include a LAN connectivity policy in a service profile, all existing vNIC configuration is erased, and when you include a SAN connectivity policy, all existing vHBA configuration in that service profile is erased.

Creating a LAN Connectivity 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 **LAN Connectivity Policies** and choose **Create LAN Connectivity Policy**.
- Step 5** In the **Create LAN Connectivity Policy** dialog box, enter a name and optional description.
- Step 6** Do one of the following:
- To add vNICs to the LAN connectivity policy, continue with Step 7.
 - To add iSCSI vNICs to the LAN connectivity policy and use iSCSI boot with the server, continue with Step 8.
- Step 7** To add vNICs, click **Add** next to the plus sign and complete the following fields in the **Create vNIC** dialog box:
- a) In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box to use an existing vNIC template.
You can also create a MAC pool from this area.
 - b) Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.
- Note** Cisco recommends using the native VLAN 1 setting to prevent traffic interruptions if using the Cisco Nexus 1000V Series Switches because changing the native VLAN 1 setting on a vNIC causes the port to turn on and off. You can only change the native VLAN setting on a Virtual Private Cloud (VPC) secondary port, and then change the primary port on the VPC.
- c) In the **Operational Parameters** area, choose a **Stats Threshold Policy**.
 - d) In the **Adapter Performance Profile** area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.
 - e) In the **Connection Policies** area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.
 - f) Click **OK**.
- Step 8** If you want to use iSCSI boot with the server, click the down arrows to expand the **Add iSCSI vNICs** bar and do the following:
- a) Click **Add** on the table icon bar.

- b) In the **Create iSCSI vNIC** dialog box, enter the **Name** and choose the **Overlay vNIC, iSCSI Adapter Policy**, and **VLAN**.

You can also create an iSCSI adapter policy from this area.

Note For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC.

For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

- c) In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, select **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.

Important If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.

- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.
- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories. **Domain Pools** are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

- d) (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.

For more information, see [Creating a MAC Pool](#).

- e) Click **OK**.

Step 9 After you have created all the vNICs or iSCSI vNICs you need for the policy, click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Creating a vNIC for a LAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies > *Organization_Name***.
 - Step 3** Expand the **LAN Connectivity Policies** node.
 - Step 4** Choose the policy to which you want to add a vNIC.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** On the icon bar of the vNICs table, click **Add**.
 - Step 7** In the **Create vNIC** dialog box, enter the name, select a **MAC Address Assignment**, and check the **Use vNIC Template** check box if you want to use an existing vNIC template.
You can also create a MAC pool from this area.
 - Step 8** Choose the **Fabric ID**, select the **VLANs** that you want to use, enter the **MTU**, and choose a **Pin Group**.
You can also create a VLAN and a LAN pin group from this area.
 - Step 9** In the **Operational Parameters** area, choose a **Stats Threshold Policy**.
 - Step 10** In the Adapter Performance Profile area, choose an **Adapter Policy**, **QoS Policy**, and a **Network Control Policy**.
You can also create an Ethernet adapter policy, QoS policy, and network control policy from this area.
 - Step 11** In the Connection Policies area, choose the **Dynamic vNIC**, **usNIC** or **VMQ** radio button, then choose the corresponding policy.
You can also create a dynamic vNIC, usNIC, or VMQ connection policy from this area.
 - Step 12** Click **OK**.
 - Step 13** Click **Save Changes**.
-

Deleting a vNIC from a LAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **LAN > Policies > *Organization_Name***.
- Step 3** Expand the **LAN Connectivity Policies** node.
- Step 4** Select the policy from which you want to delete the vNIC.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the vNICs table, do the following:
 - a) Click the vNIC you want to delete.

b) On the icon bar, click **Delete**.

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

Step 8 Click **Save Changes**.

Creating an iSCSI vNIC for a LAN Connectivity Policy

Procedure

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

Step 2 Expand **LAN > Policies > *Organization_Name***.

Step 3 Expand the **LAN Connectivity Policies** node.

Step 4 Choose the policy to which you want to add an iSCSI vNIC.

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

Step 6 On the icon bar of the **Add iSCSI vNICs** table, click **Add**.

Step 7 In the **Create iSCSI vNIC** dialog box, complete the following fields:

Name	Description
Name field	The name of the iSCSI vNIC. 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.
Overlay vNIC drop-down list	The LAN vNIC associated with this iSCSI vNIC, if any.
iSCSI Adapter Policy drop-down list	The iSCSI adapter policy associated with this iSCSI vNIC, if any.
Create iSCSI Adapter Policy link	Click this link to create a new iSCSI adapter policy that will be available to all iSCSI vNICs.
VLAN drop-down list	The virtual LAN associated with this iSCSI vNIC. The default VLAN is default . Note For the Cisco UCS M81KR Virtual Interface Card and the Cisco UCS VIC-1240 Virtual Interface Card, the VLAN that you specify must be the same as the native VLAN on the overlay vNIC. For the Cisco UCS M51KR-B Broadcom BCM57711 Adapter, the VLAN that you specify can be any VLAN assigned to the overlay vNIC.

Step 8 In the **MAC Address Assignment** drop-down list in the **iSCSI MAC Address** area, choose one of the following:

- Leave the MAC address unassigned, select **Select (None used by default)**. Select this option if the server that will be associated with this service profile contains a Cisco UCS M81KR Virtual Interface Card adapter or a Cisco UCS VIC-1240 Virtual Interface Card.
- **Important** If the server that will be associated with this service profile contains a Cisco UCS NIC M51KR-B adapter, you must specify a MAC address.
- A specific MAC address, select **00:25:B5:XX:XX:XX** and enter the address in the **MAC Address** field. To verify that this address is available, click the corresponding link.
- A MAC address from a pool, select the pool name from the list. Each pool name is followed by a pair of numbers in parentheses. The first number is the number of available MAC addresses in the pool and the second is the total number of MAC addresses in the pool.

If this Cisco UCS domain is registered with Cisco UCS Central, there might be two pool categories.

Domain Pools are defined locally in the Cisco UCS domain and **Global Pools** are defined in Cisco UCS Central.

Step 9 (Optional) If you want to create a MAC pool that will be available to all service profiles, click **Create MAC Pool** and complete the fields in the **Create MAC Pool** wizard.

For more information, see [Creating a MAC Pool](#).

Step 10 Click **OK**.

Step 11 Click **Save Changes**.

Deleting an iSCSI vNIC from a LAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
- Step 2** Expand **LAN > Policies > Organization_Name**.
- Step 3** Expand the **LAN Connectivity Policies** node.
- Step 4** Chose the policy from which you want to delete the iSCSI vNIC.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **Add iSCSI vNICs** table, do the following:
 - a) Click the iSCSI vNIC that you want to delete.
 - b) On the icon bar, click **Delete**.
- Step 7** If a confirmation dialog box displays, click **Yes**.
- Step 8** Click **Save Changes**.

Deleting a LAN Connectivity Policy

If you delete a LAN connectivity policy that is included in a service profile, it also deletes all vNICs and iSCSI vNICs from that service profile, and disrupt LAN data traffic for the server associated with the service profile.

Procedure

-
- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies > *Organization_Name***.
 - Step 3** Expand the **LAN Connectivity 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**.
-

Configuring Network Control Policies

Network Control Policy

This policy configures the network control settings for the Cisco UCS domain, including the following:

- Whether the Cisco Discovery Protocol (CDP) is enabled or disabled
- How the virtual interface (VIF) behaves if no uplink port is available in end-host mode
- The action that Cisco UCS Manager takes on the remote Ethernet interface, vEthernet interface , or vFibre Channel interface when the associated border port fails
- Whether the server can use different MAC addresses when sending packets to the fabric interconnect
- Whether MAC registration occurs on a per-VNIC basis or for all VLANs

Action on Uplink Fail

By default, the **Action on Uplink Fail** property in the network control policy is configured with a value of link-down. For adapters such as the Cisco UCS M81KR Virtual Interface Card, this default behavior directs Cisco UCS Manager to bring the vEthernet or vFibre Channel interface down if the associated border port fails. For Cisco UCS systems using a non-VM-FEX capable converged network adapter that supports both Ethernet and FCoE traffic, such as Cisco UCS CNA M72KR-Q and the Cisco UCS CNA M72KR-E, this default behavior directs Cisco UCS Manager to bring the remote Ethernet interface down if the associated border port fails. In this scenario, any vFibre Channel interfaces that are bound to the remote Ethernet interface are brought down as well.

**Note**

If your implementation includes those types of non-VM-FEX capable converged network adapters mentioned in this section and the adapter is expected to handle both Ethernet and FCoE traffic, we recommend that you configure the **Action on Uplink Fail** property with a value of warning. Note that this configuration might result in an Ethernet teaming driver not being able to detect a link failure when the border port goes down.

MAC Registration Mode

MAC addresses are installed only on the native VLAN by default, which maximizes the VLAN port count in most implementations.

**Note**

If a trunking driver is being run on the host and the interface is in promiscuous mode, we recommend that you set the MAC Registration Mode to All VLANs.

NIC Teaming and Port Security

NIC teaming is a grouping together of network adapters to build in redundancy, and is enabled on the host. This teaming or bonding facilitates various functionalities, including load balancing across links and failover. When NIC teaming is enabled and events such as failover or reconfiguration take place, MAC address conflicts and movement may happen.

Port security, which is enabled on the fabric interconnect side, prevents MAC address movement and deletion. Therefore, you must not enable port security and NIC teaming together.

Configuring Link Layer Discovery Protocol for Fabric Interconnect vEthernet Interfaces

Cisco UCS Manager Release 2.2.4 allows you to enable and disable LLDP on a vEthernet interface. You can also retrieve information about these LAN uplink neighbors. This information is useful while learning the topology of the LAN connected to the UCS system and while diagnosing any network connectivity issues from the Fabric Interconnect (FI). The FI of a UCS system is connected to LAN uplink switches for LAN connectivity and to SAN uplink switches for storage connectivity. When using Cisco UCS with Cisco Application Centric Infrastructure (ACI), LAN uplinks of the FI are connected to ACI leaf nodes. Enabling LLDP on a vEthernet interface will help the Application Policy Infrastructure Controller (APIC) to identify the servers connected to the FI by using vCenter.

To permit the discovery of devices in a network, support for Link Layer Discovery Protocol (LLDP), a vendor-neutral device discovery protocol that is defined in the IEEE 802.1ab standard, is introduced. LLDP is a one-way protocol that allows network devices to advertise information about themselves to other devices on the network. LLDP transmits information about the capabilities and current status of a device and its interfaces. LLDP devices use the protocol to solicit information only from other LLDP devices.

You can enable or disable LLDP on a vEthernet interface based on the Network Control Policy (NCP) that is applied on the vNIC in the service profile.

Creating a Network Control Policy

MAC address-based port security for Emulex converged Network Adapters (N20-AE0102) is not supported. When MAC address-based port security is enabled, the fabric interconnect restricts traffic to packets that contain the MAC address that it first learns. This is either the source MAC address used in the FCoE Initialization Protocol packet, or the MAC address in an ethernet packet, whichever is sent first by the adaptor. This configuration can result in either FCoE or Ethernet packets being dropped.

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 **Network Control Policies** node and select **Create Network Control Policy**.
- Step 5** In the **Create Network Control Policy** dialog box, complete the required fields.
- Step 6** In the LLDP area, do the following:
- To enable the transmission of LLDP packets on an interface, click **Enabled** in the **Transmit** field.
 - To enable the reception of LLDP packets on an interface, click **Enabled** in the **Receive** field.
- Step 7** In the **MAC Security** area, do the following to determine whether the server can use different MAC addresses when sending packets to the fabric interconnect:
- Click the **Expand** icon to expand the area and display the radio buttons.
 - Click one of the following radio buttons to determine whether forged MAC addresses are allowed or denied when packets are sent from the server to the fabric interconnect:
 - **Allow**— All server packets are accepted by the fabric interconnect, regardless of the MAC address associated with the packets.
 - **Deny**— After the first packet has been sent to the fabric interconnect, all other packets must use the same MAC address or they will be silently rejected by the fabric interconnect. In effect, this option enables port security for the associated vNIC.
- If you plan to install VMware ESX on the associated server, you must configure the **MAC Security** to **allow** for the network control policy applied to the default vNIC. If you do not configure **MAC Security** for **allow**, the ESX installation may fail because the MAC security permits only one MAC address while the installation process requires more than one MAC address.
- Step 8** Click **OK**.
-

Deleting a Network Control Policy

Procedure

-
- | | |
|---------------|--|
| Step 1 | In the Navigation pane, click LAN . |
| Step 2 | Expand LAN > Policies > <i>Organization_Name</i> . |
| Step 3 | Expand the Network Control Policies node. |
| Step 4 | Right-click the policy you want to delete and select Delete . |
| Step 5 | If a confirmation dialog box displays, click Yes . |
-

Configuring Multicast Policies

Multicast Policy

This policy is used to configure Internet Group Management Protocol (IGMP) snooping and IGMP querier. IGMP Snooping dynamically determines hosts in a VLAN that should be included in particular multicast transmissions. You can create, modify, and delete a multicast policy that can be associated to one or more VLANs. When a multicast policy is modified, all VLANs associated with that multicast policy are re-processed to apply the changes. For private VLANs, you can set a multicast policy for primary VLANs but not for their associated isolated VLANs due to a Cisco NX-OS forwarding implementation.

By default, IGMP snooping is enabled and IGMP querier is disabled. When IGMP snooping is enabled, the fabric interconnects send the IGMP queries only to the hosts. They do not send IGMP queries to the upstream network. To send IGMP queries to the upstream, do one of the following:

- Configure IGMP querier on the upstream fabric interconnect with IGMP snooping enabled
- Disable IGMP snooping on the upstream fabric interconnect
- Change the fabric interconnects to switch mode

The following limitations and guidelines apply to multicast policies:

- On a 6200 series fabric interconnect, user-defined multicast policies can also be assigned along with the default multicast policy.
- Only the default multicast policy is allowed for a global VLAN.
- If a Cisco UCS domain includes 6300 and 6200 series fabric interconnects, any multicast policy can be assigned.
- We highly recommend you use the same IGMP snooping state on the fabric interconnects and the associated LAN switches. For example, if IGMP snooping is disabled on the fabric interconnects, it should be disabled on any associated LAN switches as well.

Creating a Multicast Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies**.
 - Step 3** Expand the **root** node.
 - Step 4** Right-click the **Multicast Policies** node and select **Create Multicast Policy**.
 - Step 5** In the **Create Multicast Policy** dialog box, specify the name and IGMP snooping information.
 - Step 6** Click **OK**.
-

Modifying a Multicast Policy

This procedure describes how to change the IGMP snooping state and the IGMP snooping querier state of an existing multicast policy.



Note

You cannot change the name of the multicast policy once it has been created.

Procedure

-
- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies**.
 - Step 3** Expand the **root** node.
 - Step 4** Click the policy that you want to modify.
 - Step 5** In the work pane, edit the fields as needed.
 - Step 6** Click **Save Changes**.
-

Deleting a Multicast Policy



Note

If you assigned a non-default (user-defined) multicast policy to a VLAN and then delete that multicast policy, the associated VLAN inherits the multicast policy settings from the default multicast policy until the deleted policy is re-created.

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies**.
 - Step 3** Expand the **root** node.
 - Step 4** Right-click the **Multicast Policies** node and select **Delete Multicast Policy**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Configuring LACP policies

LACP Policy

Link Aggregation combines multiple network connections in parallel to increase throughput and to provide redundancy. Link aggregation control protocol (LACP) provides additional benefits for these link aggregation groups. Cisco UCS Manager enables you to configure LACP properties using LACP policy.

You can configure the following for a lacp policy:

- **Suspended-individual:** If you do not configure the ports on an upstream switch for lacp, the fabric interconnects treat all ports as uplink Ethernet ports to forward packets. You can place the lacp port in suspended state to avoid loops. When you set suspend-individual on a port-channel with LACP, if a port that is part of the port-channel does not receive PDUs from the peer port, it will go into suspended state.
- **Timer values:** You can configure rate-fast or rate-normal. In rate-fast configuration, the port is expected to receive 1 PDU every 1 second from the peer port. The time out for this is 3 seconds. In rate-normal configuration, the port is expected to receive 1 PDU every 30 seconds. The timeout for this is 90 seconds.

System creates a default LACP policy at system start up. You can modify this policy or create a new policy. You can also apply one LACP policy to multiple port-channels.

Creating a LACP 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** In the **Work Pane**, click **LACP Policies** tab, and click the + sign.
 - Step 5** In the **Create LACP Policy** dialog box, fill in the required fields.
 - Step 6** Click **OK**.
-

Modifying a LACP 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** In the **Work Pane**, **LACP Policies** tab, and click on the policy you want to edit.
 - Step 5** Click the **Properties** icon on the right.
 - Step 6** In the **Properties** dialog box, make the required changes and click **Apply**.
 - Step 7** Click **OK**.
-

Configuring UDLD Link Policies

Understanding UDLD

UniDirectional Link Detection (UDLD) is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it marks the link as unidirectional. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, the Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected interfaces on fiber-optic connections. In aggressive mode,

UDLD can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected interfaces on fiber-optic links.

In normal mode, UDLD detects a unidirectional link when fiber strands in a fiber-optic interface are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the interfaces are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In case, the logical link is considered undetermined, and UDLD does not disable the interface. When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected and autonegotiation is active, the link does not stay up because the Layer 1 mechanisms did not detect a physical problem with the link. In this case, UDLD does not take any action, and the logical link is considered undetermined.

UDLD aggressive mode is disabled by default. Configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. With UDLD aggressive mode enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD packets, UDLD tries to reestablish the connection with the neighbor and administratively shuts down the affected port. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of the following problems exists:

- On fiber-optic or twisted-pair links, one of the interfaces cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the interfaces is down while the other is up.
- One of the fiber strands in the cable is disconnected.

Methods to Detect Unidirectional Links

UDLD operates by using two mechanisms:

- Neighbor database maintenance

UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors. When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

UDLD clears all existing cache entries for the interfaces affected by the configuration change whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply.

If the detection window ends and no valid reply message is received, the link might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and might not be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the interface is shut down.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of a port have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the port if, after the fast train of messages, the link state is still undetermined.

UDLD Configuration Guidelines

The following guidelines and recommendations apply when you configure UDLD:

- A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.
- UDLD should be enabled only on interfaces that are connected to UDLD capable devices. The following interface types are supported:
 - Ethernet uplink
 - FCoE uplink
 - Ethernet uplink port channel member
 - FCoE uplink port channel member

Creating a Link Profile

Procedure

-
- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies > LAN Cloud**.
 - Step 3** Right-click the **Link Profile** node and choose **Create Link Profile**.
 - Step 4** In the **Create Link Profile** dialog box, specify the name and the UDLD link policy.
 - Step 5** Click **OK**.
-

Creating a UDLD Link Policy

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies > LAN Cloud**.
 - Step 3** Right-click the **UDLD Link Policy** node and choose **Create UDLD Link Policy**.
 - Step 4** In the **Create UDLD Link Policy** dialog box, specify the name, admin state, and mode.
 - Step 5** Click **OK**.
-

Modifying the UDLD System Settings

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > Policies > LAN Cloud**.
 - Step 3** On the **LAN** tab, expand **LAN > Policies > root**.
 - Step 4** Expand the **Link Protocol Policy** node and click **UDLD System Settings**.
 - Step 5** In the **Work** pane, click the **General** tab.
 - Step 6** In the **Properties** area, modify the fields as needed.
 - Step 7** Click **Save Changes**.
-

Assigning a Link Profile to a Port Channel Ethernet Interface

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** Expand **LAN > LAN Cloud > Fabric > Port Channels**.
 - Step 3** Expand the port channel node and click the Eth Interface where you want to assign a link profile.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Properties** area, choose the link profile that you want to assign.
 - Step 6** Click **Save Changes**.
-

Assigning a Link Profile to an Uplink Ethernet Interface

Procedure

- Step 1** In the **Navigation** pane, click **LAN**.
 - Step 2** On the **LAN** tab, expand **LAN > LAN Cloud > Fabric > Uplink Eth Interface**
 - Step 3** Click the Eth Interface where you want to assign a link profile.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Properties** area, choose the link profile that you want to assign.
 - Step 6** Click **Save Changes**.
-

Assigning a Link Profile to a Port Channel FCoE Interface

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** On the **SAN** tab, expand **SAN > SAN Cloud > Fabric > FCoE Port Channels**
 - Step 3** Expand the FCoE port channel node and click the FCoE Interface where you want to assign a link profile.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Properties** area, choose the link profile that you want to assign.
 - Step 6** Click **Save Changes**.
-

Assigning a Link Profile to an Uplink FCoE Interface

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** On the **SAN** tab, expand **SAN > SAN Cloud > Fabric > Uplink FC Interfaces**
 - Step 3** Click the FCoE interface where you want to assign a link profile.
 - Step 4** In the **Work** pane, click the **General** tab.
 - Step 5** In the **Properties** area, choose the link profile that you want to assign.
 - Step 6** Click **Save Changes**.
-

Configuring VMQ Connection Policies

VMQ Connection Policy

Cisco UCS Manager enables you to configure VMQ connection policy for a vNIC. VMQ provides improved network performance to the entire management operating system. Configuring a VMQ vNIC connection policy involves the following:

- Create a VMQ connection policy
- Create a static vNIC in a service profile
- Apply the VMQ connection policy to the vNIC

If you want to configure the VMQ vNIC on a service profile for a server, at least one adapter in the server must support VMQ. Make sure the servers have at least one the following adapters installed:

- UCS-VIC-M82-8P
- UCSB-MLOM-40G-01
- UCSC-PCIE-CSC-02

The following are the supported Operating Systems for VMQ:

- Windows 2012
- Windows 2012R2

You can apply only any one of the vNIC connection policies on a service profile at any one time. Make sure to select one of the three options such as Dynamic, usNIC or VMQ connection policy for the vNIC. When a VMQ vNIC is configured on service profile, make sure you have the following settings:

- Select SRIOV in the BIOS policy.
- Select Windows in the Adapter policy.

Creating a VMQ Connection Policy

Before you create a VMQ connection policy, consider the following:

- VMQ Tuning on the Windows Server — When an adapter is placed on a virtual switch, running the **Get-NetAdapterVmq** cmdlet displays True for VMQ. For more information on NIC teaming see [Performance Tuning for Hyper-V Servers](#).
- Virtual machine level — By default, VMQ is enabled on all newly deployed VMs. VMQ can be enabled or disabled on existing VMs.
- Microsoft SCVMM — VMQ must be enabled on the port profile. If not, you will not be able to successfully create the virtual switch in SCVMM.

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 **VMQ Connection Policies** node and select **Create VMQ Connection Policy**.
- Step 5** In the **Create VMQ Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	The VMQ connection policy name.
Description field	The description of the VMQ connection policy.
Number of VMQs field	The number of VMQs per adapter must be one more than the maximum number of VM NICs. Note Make sure that the total number of synthetic NICs present on the VMs is either equal to or greater than the number of VMs.
Number of Interrupts field	The number of CPU threads or logical processors available in the server. Note You cannot set this value to be more than the maximum number of available CPUs.

- Step 6** Click **OK**.

Assigning Virtualization Preference to a vNIC

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** On the **Servers** tab, expand **Servers > target service profile > root > vNICs**.
- Step 3** Click on the vNIC name to display properties on the work pane.
- Step 4** In the **Connection Policies** section, select the radio button for **VMQ** and select the **VMQ Connection Policy** from the drop down.
In the **Properties** area **Virtualization Preference** for this vNIC changes to **VMQ**.

Enabling VMQ and NVGRE Offloading on the same vNIC

Perform the tasks in the table below to enable VMQ and NVGRE offloading on the same vNIC.


Note

Currently, VMQ is not supported along with VXLAN on the same vNIC.

Task	Description	See
Enable normal NVGRE offloading	Perform this task by setting the corresponding flags in the adapter profile which is associated with the given vNIC. Note The Transmit checksum offload and TSO must be enabled for the NVGRE offloading to be effective.	Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE, on page 17 Applying an NVGRE Adapter Policy to a vNIC, on page 40
Enable VMQ	Perform this task by setting the appropriate connection policy when you add a vNIC to the service profile.	Creating a VMQ Connection Policy, on page 37 Assigning Virtualization Preference to a vNIC, on page 38

Configuring an Ethernet Adapter Policy to Enable Stateless Offloads with NVGRE

Cisco UCS Manager supports stateless offloads with NVGRE only with Cisco UCS VIC 1340 and/or Cisco UCS VIC 1380 adapters that are installed on servers running Windows Server 2012 R2 operating systems. Stateless offloads with NVGRE cannot be used with NetFlow, usNIC, or VM-FEX.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > 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 **Adapter Policies** and choose **Create Ethernet Adapter Policy**.
 - a) In the **Resources** area, set the following options:
 - Transmit Queues = 1
 - Receive Queues = n (up to 8)
 - Completion Queues = # of Transmit Queues + # of Receive Queues

- Interrupts = # Completion Queues + 2

b) In the **Options** area, set the following options:

- Network Virtualization using Generic Routing Encapsulation = Enabled
- Interrupt Mode = Msi-X

For more information on creating an Ethernet adapter policy, see [Creating an Ethernet Adapter Policy](#), on page 12.

Step 5 Click **OK** to create the Ethernet adapter policy.

Step 6 Install an eNIC driver Version 3.0.0.8 or later.

For more information, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows/b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.html.

Step 7 Reboot the server.

Applying an NVGRE Adapter Policy to a vNIC

Procedure

Step 1 In the **Navigation** pane, click the **Servers** tab.

Step 2 On the **Servers** tab, expand **Servers > Target Service Profile > root > vNICS**

Step 3 Click on the vNIC name to display properties in the work pane.

Step 4 In the **Policies** section, select the NVGRE policy from **Adapter Policy** drop-down list.

Step 5 Click **Save Changes** to apply the policy to the vNIC.

Creating a VMQ Connection Policy

Before you create a VMQ connection policy, consider the following:

- VMQ Tuning on the Windows Server — When an adapter is placed on a virtual switch, running the **Get-NetAdapterVmq** cmdlet displays True for VMQ. For more information on NIC teaming see [Performance Tuning for Hyper-V Servers](#).
- Virtual machine level — By default, VMQ is enabled on all newly deployed VMs. VMQ can be enabled or disabled on existing VMs.
- Microsoft SCVMM — VMQ must be enabled on the port profile. If not, you will not be able to successfully create the virtual switch in SCVMM.

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 **VMQ Connection Policies** node and select **Create VMQ Connection Policy**.
- Step 5** In the **Create VMQ Connection Policy** dialog box, complete the following fields:

Name	Description
Name field	The VMQ connection policy name.
Description field	The description of the VMQ connection policy.
Number of VMQs field	The number of VMQs per adapter must be one more than the maximum number of VM NICs. Note Make sure that the total number of synthetic NICs present on the VMs is either equal to or greater than the number of VMs.
Number of Interrupts field	The number of CPU threads or logical processors available in the server. Note You cannot set this value to be more than the maximum number of available CPUs.

- Step 6** Click **OK**.

Assigning Virtualization Preference to a vNIC

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** On the **Servers** tab, expand **Servers > target service profile > root > vNICs**.
- Step 3** Click on the vNIC name to display properties on the work pane.
- Step 4** In the **Connection Policies** section, select the radio button for **VMQ** and select the **VMQ Connection Policy** from the drop down.
In the **Properties** area **Virtualization Preference** for this vNIC changes to **VMQ**.

NetQueue

Information About NetQueue

NetQueue improves traffic performance by providing a network adapter with multiple receive queues. These queues allow the data interrupt processing that is associated with individual virtual machines to be grouped.

**Note**

NetQueue is supported on servers running VMware ESXi operating systems.

Configuring NetQueue

Procedure

Step 1 Create a Virtual Machine Queue (VMQ) connection policy.

Step 2 Configure NetQueues in a service profile by selecting the VMQ connection policy.
Use the following when you are configuring NetQueue:

- The default ring size is rx512, tx256
- The interrupt count on each VNIC is VMQ count x 2 + 2

Note The number of interrupts depends on the number of NetQueues enabled.

- The driver supports up to 16 NetQueues per port for standard frame configurations.

Note VMware recommends that you use up to eight NetQueues per port for standard frame configurations.

- NetQueue should be enabled only on MSIX systems.
- You should disable NetQueue on 1 GB NICs.

Step 3 Enable the MSIX mode in the adapter policy for NetQueue.

Step 4 Associate the service profile with the server.