



Managing Network Adapters

This chapter includes the following sections:

- [Overview of the Cisco UCS C-Series Network Adapters, on page 1](#)
- [Configuring Network Adapter Properties, on page 3](#)
- [Managing vHBAs, on page 11](#)
- [Managing vNICs, on page 25](#)
- [Backing Up and Restoring the Adapter Configuration, on page 50](#)
- [Resetting the Adapter, on page 53](#)

Overview of the Cisco UCS C-Series Network Adapters



Note The procedures in this chapter are available only when a Cisco UCS C-Series network adapter is installed in the chassis.

A Cisco UCS C-Series network adapter can be installed to provide options for I/O consolidation and virtualization support. The following adapters are available:

- Cisco UCS VIC 15238 Virtual Interface Card
- Cisco UCS VIC 15428 Virtual Interface Card
- Cisco UCS VIC 1497 Virtual Interface Card
- Cisco UCS VIC 1495 Virtual Interface Card
- Cisco UCS VIC 1477 Virtual Interface Card
- Cisco UCS VIC 1467 Virtual Interface Card
- Cisco UCS VIC 1457 Virtual Interface Card
- Cisco UCS VIC 1455 Virtual Interface Card
- Cisco UCS VIC 1387 Virtual Interface Card
- Cisco UCS VIC 1385 Virtual Interface Card
- Cisco UCS VIC 1227T Virtual Interface Card

- Cisco UCS VIC 1225 Virtual Interface Card



Note You must have same generation VIC cards on a server. For example, you cannot have a combination of 3rd generation and 4th generation VIC cards on a single server.

The interactive *UCS Hardware and Software Interoperability Utility* lets you view the supported components and configurations for a selected server model and software release. The utility is available at the following URL: <http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>

Cisco UCS VIC 15428 Virtual Interface Card

The Cisco VIC 15428 is a quad-port Small Form-Factor Pluggable (SFP+/SFP28/SFP56) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack servers. The card supports 10/25/50-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Cisco UCS VIC 1497 Virtual Interface Card

The Cisco VIC 1497 is a dual-port Small Form-Factor (QSFP28) mLOM card designed for the M5 generation of Cisco UCS C-Series Rack Servers. The card supports 40/100-Gbps Ethernet and FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs and HBAs.

Cisco UCS VIC 1495 Virtual Interface Card

The Cisco UCS VIC 1495 is a dual-port Small Form-Factor (QSFP28) PCIe card designed for the M5 generation of Cisco UCS C-Series Rack Servers. The card supports 40/100-Gbps Ethernet and FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs and HBAs.

Cisco UCS VIC 1477 Virtual Interface Card

The Cisco VIC 1477 is a dual-port Quad Small Form-Factor (QSFP28) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack Servers. The card supports 40/100-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs or HBAs.

Cisco UCS VIC 1467 Virtual Interface Card

The Cisco UCS VIC 1467 is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBA.

Cisco UCS VIC 1457 Virtual Interface Card

The Cisco UCS VIC 1457 is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for M5 generation of Cisco UCS C-Series rack servers. The card supports 10/25-Gbps Ethernet or FCoE. It incorporates Cisco's next-generation CNA technology and offers a comprehensive feature set, providing investment protection for future feature software releases. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs and HBAs.

Cisco UCS VIC 1455 Virtual Interface Card

The Cisco UCS VIC 1455 is a quad-port Small Form-Factor Pluggable (SFP28) half-height PCIe card designed for M5 generation of Cisco UCS C-Series rack servers. The card supports 10/25-Gbps Ethernet or FCoE. It incorporates Cisco's next-generation CNA technology and offers a comprehensive feature set, providing investment protection for future feature software releases. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as NICs and HBAs.

Cisco UCS VIC 1387 Virtual Interface Card

The Cisco UCS VIC 1387 Virtual Interface Card is a dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP) 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable half-height PCI Express (PCIe) card designed exclusively for Cisco UCS C-Series Rack Servers. It incorporates Cisco's next-generation converged network adapter (CNA) technology, with a comprehensive feature set, providing investment protection for future feature software releases.

Cisco UCS VIC 1385 Virtual Interface Card

The Cisco UCS VIC 1385 Virtual Interface Card is a dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP) 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable half-height PCI Express (PCIe) card designed exclusively for Cisco UCS C-Series Rack Servers. It incorporates Cisco's next-generation converged network adapter (CNA) technology, with a comprehensive feature set, providing investment protection for future feature software releases.

Cisco UCS VIC 1227T Virtual Interface Card

The Cisco UCS VIC 1227T Virtual Interface Card is a dual-port 10GBASE-T (RJ-45) 10-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter designed exclusively for Cisco UCS C-Series Rack Servers. New to Cisco rack servers, the mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing Fibre Channel connectivity over low-cost twisted pair cabling with a bit error rate (BER) of 10 to 15 up to 30 meters and investment protection for future feature releases.

Cisco UCS VIC 1225 Virtual Interface Card

The Cisco UCS VIC 1225 Virtual Interface Card is a high-performance, converged network adapter that provides acceleration for the various new operational modes introduced by server virtualization. It brings superior flexibility, performance, and bandwidth to the new generation of Cisco UCS C-Series Rack-Mount Servers.

Configuring Network Adapter Properties

Before you begin

- You must log in as a user with admin privileges to perform this task.
- The server must be powered on, or the properties will not display.

Step 1 In the **Navigation** pane, click the **Networking** menu.

Step 2 In the **Networking** menu, select the adapter card that you want to view.

Step 3 In the **Adapter Card Properties** area under the **General** tab, review the following information:

Name	Description
PCI Slot field	The PCI slot in which the adapter is installed.
Vendor field	The vendor for the adapter.
Product Name field	The product name for the adapter.
Product ID field	The product ID for the adapter.
Serial Number field	The serial number for the adapter.
Version ID field	The version ID for the adapter.
PCI Link field	The server to which the PCIe link is established.
Hardware Revision field	The hardware revision for the adapter.
Cisco IMC Management Enabled field	If this field displays yes , then the adapter is functioning in Cisco Card Mode and passing Cisco IMC management traffic through to the server Cisco IMC.
Configuration Pending field	If this field displays yes , the adapter configuration has changed in Cisco IMC but these changes have not been communicated to the host operating system. To activate the changes, an administrator must reboot the adapter.
ISCSI Boot Capable field	Whether iSCSI boot is supported on the adapter.
CDN Capable field	Whether CDN is supported on the adapter.
usNIC Capable field	Whether the adapter and the firmware running on the adapter support the usNIC.
Port Channel Capable field	Indicates whether Port Channel is supported on the adapter. Note This option is available only on some of the adapters and servers.
Description field	A user-defined description for the adapter. You can enter between 1 and 63 characters.
Enable FIP Mode check box	If checked, then FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards. Note We recommend that you use this option only when explicitly directed to do so by a technical support representative.

Name	Description
<p>Enable LLDP check box</p>	<p>Note For LLDP change to be effective, it is required that you reboot the server.</p> <p>In case of S3260 chassis with two nodes, ensure to reboot the secondary node after making LLDP changes in the primary node.</p> <p>If checked, then Link Layer Discovery Protocol (LLDP) enables all the Data Center Bridging Capability Exchange protocol (DCBX) functionality, which includes FCoE, priority based flow control.</p> <p>By default, LLDP option is enabled.</p> <p>Note We recommend that you do not disable LLDP option, as it disables all the DCBX functionality.</p>
<p>Enable VNTAG Mode check box</p>	<p>If VNTAG mode is enabled:</p> <ul style="list-style-type: none"> • vNICs and vHBAs can be assigned to a specific channel. • vNICs and vHBAs can be associated to a port profile. • vNICs can fail over to another vNIC if there are communication problems.
<p>Port Channel check box</p>	<p>This option is enabled by default.</p> <p>When Port channel is enabled, two vNICs and two vHBAs are available for use on the adapter card.</p> <p>When disabled, four vNICs and four vHBAs are available for use on the adapter card.</p> <p>Note This option is available only on some of the adapters and servers.</p>

Name	Description
Physical NIC Mode check box	<p>This option is disabled by default.</p> <p>When Physical NIC Mode is enabled, up-link ports of the VIC are set to pass-through mode. This allows the host to transmit packets without any modification. VIC ASIC does not rewrite the VLAN tag of the packets based on the VLAN and CoS settings for the vNIC.</p> <p>Note</p> <ul style="list-style-type: none"> • This option is available for Cisco UCS VIC 14xx series and 15xxx series adapters. • For the VIC configuration changes to be effective, you must reboot the host. • This option cannot be enabled on an adapter that has: <ul style="list-style-type: none"> • Port Channel mode enabled • VNTAG mode enabled • LLDP enabled • FIP mode enabled • Cisco IMC Management Enabled value set to Yes • Multiple user created vNICs <p>When Physical NIC Mode is enabled, the following message is displayed in a pop-up window:</p> <p>After physical nic-mode mode switch, vNIC configurations will be lost and new default vNICs will be created.</p> <p>Click OK.</p>

Step 4 In the **Firmware** area, review the following information:

Name	Description
Running Version field	The firmware version that is currently active.

Name	Description
Backup Version field	The alternate firmware version installed on the adapter, if any. The backup version is not currently running. To activate it, administrators can click Activate Firmware in the Actions area. Note When you install new firmware on the adapter, any existing backup version is deleted and the new firmware becomes the backup version. You must manually activate the new firmware if you want the adapter to run the new version.
Startup Version field	The firmware version that will become active the next time the adapter is rebooted.
Bootloader Version field	The bootloader version associated with the adapter card.
Status field	The status of the last firmware activation that was performed on this adapter. Note The status is reset each time the adapter is rebooted.

Step 5 Click the **External Ethernet Interfaces** link to review the following information:

Note **External Ethernet Interfaces** opens in a different tab.

Name	Description
Port column	The uplink port ID.
Admin Speed column	The data transfer rate for the port. This can be one of the following: <ul style="list-style-type: none"> • 40 Gbps • 4 x 10 Gbps • Auto
Admin Link Training column	Indicates if admin link training is enabled on the port. Select any of the below options for Admin Link Training: <ul style="list-style-type: none"> • Auto • Off • On Admin Link Training is set to Auto , by default. Note This option is available only on some of the adapters and servers.

Name	Description
Admin FEC Mode drop-down list	<p>Admin Forward Error Correction (FEC) apply only to Cisco UCS VIC 14xx adapters at speed 25/100G and Cisco UCS VIC 15xxx adapters at speeds 25G/50G.</p> <p>Following Admin Forward Error Correction (FEC) mode options are available for configuration:</p> <ul style="list-style-type: none"> • cl108 (RS-IEEE, clause 108) • cl91-cons16 (RS-FEC, clause 91, consortium version 1.6) • cl91 (RS-FEC, clause 91, consortium version 1.5) • cl74 (FC-FEC, clause 74), 25G only • Off <p>Admin FEC Mode is set to cl91, by default.</p> <p>Note</p> <ul style="list-style-type: none"> • This option is available only on some of the adapters and servers. • Any changes in the Admin FEC Mode settings leads to the reset of the port, even when the value for Operating FEC mode remains the same.
Operating FEC Mode column	<p>The value of Operating FEC Mode is the same as Admin FEC mode with these exceptions:</p> <ul style="list-style-type: none"> • The value is Off when the speed is 10 Gbps or 40 Gbps. This is because FEC is not supported. • The value is Off for QSFP-100G-LR4-S transceiver. • The value is Off for QSFP-40/100-SRBD transceiver. <p>Note This option is available only on some of the adapters and servers.</p>

Name	Description
Oper Link Training column	<p>Oper Link Training values are fetched from the values set in the Admin Link Training drop-down list.</p> <p>Beginning from 4.2(2a), the below different settings apply only to Cisco UCS VIC 15xxx adapters and Copper cables at speeds 10G/25G/50G only.</p> <ul style="list-style-type: none"> • If Admin Link Training is set to Auto, Adapter firmware sets Oper Link Training value (AutoNeg) as on or off, depending upon the transceivers. <ul style="list-style-type: none"> • AutoNeg disabled with 25G copper • AutoNeg enabled with 50G copper • If Admin Link Training is set to on, Adapter firmware sets Oper Link Training value as on. <ul style="list-style-type: none"> • AutoNeg enabled with 25G copper • AutoNeg enabled with 50G copper • If Admin Link Training is off, Adapter firmware sets Oper Link Training as off. <ul style="list-style-type: none"> • AutoNeg disabled with 25G copper • AutoNeg disabled with 50G copper <p>Note For all non-passive copper cables, Oper Link Training (AN) mode is set to Off, irrespective of the Admin Link Training mode.</p> <p>Any changes in the Admin Link Training settings leads to the reset of the Series for that port, even if the Oper Link Training value remains the same.</p>
MAC Address column	The MAC address of the uplink port.

Name	Description
Link State column	<p>The current operational state of the uplink port. This can be one of the following:</p> <ul style="list-style-type: none"> • Fault • Link Up • Link Down • SFP ID Error • SFP Not Installed • SFP Security Check Failed • Unsupported SFP <p>Note</p> <ul style="list-style-type: none"> • A Serdes reset causes the Link State field to change from Link-Up to Link-down. <p>If the Oper Link Training setting is valid, Link-Partners determine a Link-Up or Link-down after the reset.</p> <ul style="list-style-type: none"> • You might require to refresh the Web UI several times to view Link State field change.
Encap column	<p>The mode in which adapter operates. This can be one of the following:</p> <ul style="list-style-type: none"> • CE—Classical Ethernet mode. • NIV—Network Interface Virtualization mode.
Operating Speed column	<p>The operating rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> • 10 Gbps • 25 Gbps • 40 Gbps • 50 Gbps • 4 x 10 Gbps • 100 Gbps
Connector Present column	<p>Indicated whether or not the connector is present. This can be one of the following:</p> <ul style="list-style-type: none"> • Yes—Connector is present. • No—Connector not present. <p>Note This option is only available for some adapter cards.</p>

Name	Description
Connector Supported column	<p>Indicates whether or not the connector is supported by Cisco. This can be one of the following:</p> <ul style="list-style-type: none"> • Yes—The connector is supported by Cisco. • No—The connector is not supported by Cisco. <p>If the connector is not supported then the link will not be up.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Type column	<p>The Cisco Product ID (PID) of the transceiver/cable that is present.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Vendor column	<p>The vendor for the connector.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Part Number column	<p>The Connector Vendor part number.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Part Revision column	<p>The part revision of the Connector Vendor part number.</p> <p>Note This option is only available for some adapter cards.</p>

Managing vHBAs

Guidelines for Managing vHBAs

When managing vHBAs, consider the following guidelines and restrictions:

- The Cisco UCS Virtual Interface Cards provide two vHBAs and two vNICs by default. You can create up to 14 additional vHBAs or vNICs on these adapter cards.

The Cisco UCS 1455, 1457, and 1467 Virtual Interface Cards, in non-port channel mode, provide four vHBAs and four vNICs by default. You can create up to 10 additional vHBAs or vNICs on these adapter cards in VNTAG mode.



Note If VNTAG mode is enabled for the adapter, you must assign a channel number to a vHBA when you create it.

- When using the Cisco UCS Virtual Interface Cards in an FCoE application, you must associate the vHBA with the FCoE VLAN. Follow the instructions in the **Modifying vHBA Properties** section to assign the VLAN.

- After making configuration changes, you must reboot the host for settings to take effect.

Viewing vHBA Properties

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to view.
- Step 3** In the **Adapter Card** area, click the **vHBAs** tab.
- Step 4** In the **vHBAs** pane, click **fc0** or **fc1**.
- Step 5** In the **General** area of vHBA Properties, review the information in the following fields:

Name	Description
Name field	The name of the virtual HBA. This name cannot be changed after the vHBA has been created.
Initiator WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
Initiator WWPN field	The WWPN associated with the vHBA. To let the system generate the WWPN, select AUTO . To specify a WWPN, click the second radio button and enter the WWPN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
Persistent LUN Binding check box	If checked, any LUN ID associations are retained in memory until they are manually cleared.
Uplink Port drop-down list	The uplink port associated with the vHBA. Note This value cannot be changed for the system-defined vHBAs fc0 and fc1.
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
PCI Order field	The order in which this vHBA will be used. To let the system set the order, select ANY . To specify an order, select the second radio button and enter an integer between 0 and 17.

Name	Description
vHBA Type drop-down list	<p>Note This option is available only with 14xx series and VIC 15428 adapters.</p> <p>The vHBA type used in this policy. vHBAs supporting FC and FC-NVMe can now be created on the same adapter. The vHBA type used in this policy can be one of the following:</p> <ul style="list-style-type: none"> • fc-initiator—Legacy SCSI FC vHBA initiator • fc-target—vHBA that supports SCSI FC target functionality <p>Note This option is available as a Tech Preview.</p> <ul style="list-style-type: none"> • fc-nvme-initiator—vHBA that is an FC NVMe initiator, which discovers FC NVMe targets and connects to them. • fc-nvme-target—vHBA that acts as an FC NVMe target and provides connectivity to the NVMe storage.
Class of Service field	<p>The CoS for the vHBA.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Rate Limit field	<p>The data rate limit for traffic on this vHBA, in Mbps.</p> <p>If you want this vHBA to have an unlimited data rate, select OFF. Otherwise, click the second radio button and enter an integer between 1 and 10,000.</p> <p>Note This option cannot be used in VNTAG mode.</p>
EDTOV field	<p>The error detect timeout value (EDTOV), which is the number of milliseconds to wait before the system assumes that an error has occurred.</p> <p>Enter an integer between 1,000 and 100,000. The default is 2,000 milliseconds.</p>
RATOV field	<p>The resource allocation timeout value (RATOV), which is the number of milliseconds to wait before the system assumes that a resource cannot be properly allocated.</p> <p>Enter an integer between 5,000 and 100,000. The default is 10,000 milliseconds.</p>
Max Data Field Size field	<p>The maximum size of the Fibre Channel frame payload bytes that the vHBA supports.</p> <p>Enter an integer between 256 and 2112.</p>
Channel Number field	<p>The channel number that will be assigned to this vHBA.</p> <p>Enter an integer between 1 and 1,000.</p> <p>Note VNTAG mode is required for this option.</p>

Name	Description
PCI Link	It is read-only field.
Port Profile drop-down list	<p>The port profile that should be associated with the vHBA, if any.</p> <p>This field displays the port profiles defined on the switch to which this server is connected.</p> <p>Note VNTAG mode is required for this option.</p>

Step 6 In the **Error Recovery** area, review the information in the following fields:

Name	Description
FCP Error Recovery check box	If checked, the system uses FCP Sequence Level Error Recovery protocol (FC-TAPE).
Link Down Timeout field	<p>The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost.</p> <p>Enter an integer between 0 and 240,000.</p>
Port Down I/O Retry Count field	<p>The number of times an I/O request to a port is returned because the port is busy before the system decides the port is unavailable.</p> <p>Enter an integer between 0 and 255.</p>
IO Timeout Retry field	<p>The time period till which the system waits for timeout before retrying. When a disk does not respond for I/O within the defined timeout period, the driver aborts the pending command, and resends the same I/O after the timer expires.</p> <p>Enter an integer between 1 and 59.</p>
Port Down Timeout field	<p>The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable.</p> <p>Enter an integer between 0 and 240,000.</p>

Step 7 In the **Fibre Channel Interrupt** area, review the information in the following fields:

Name	Description
Interrupt Mode drop-down list	<p>The preferred driver interrupt mode. This can be one of the following:</p> <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—MSI only. • INTx—PCI INTx interrupts.

Step 8 In the **Fibre Channel Port** area, review the information in the following fields:

Name	Description
I/O Throttle Count field	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1,024.
LUNs per Target field	The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation. Enter an integer between 1 and 4,096.
LUN Queue Depth field	The number of commands that the HBA can send or receive in a single chunk per LUN. This parameter adjusts the initial queue depth for all LUNs on the adapter. Default value is 20 for physical miniports and 250 for virtual miniports.

Step 9 In the **Fibre Channel Port FLOGI** area, review the information in the following fields:

Name	Description
FLOGI Retries field	The number of times that the system tries to log in to the fabric after the first failure. To specify an unlimited number of retries, select the INFINITE radio button. Otherwise select the second radio button and enter an integer into the corresponding field.
FLOGI Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1,000 and 255,000.

Step 10 In the **Fibre Channel Port PLOGI** area, review the information in the following fields:

Name	Description
PLOGI Retries field	The number of times that the system tries to log in to a port after the first failure. Enter an integer between 0 and 255.
PLOGI Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1,000 and 255,000.

Step 11 In the **I/O** area, review the information in the following fields:

Name	Description
CDB Transmit Queue Count field	The number of SCSI I/O queue resources the system should allocate. For Cisco UCS VIC 14xx series adapters, enter an integer between 1 and 64. For any other VIC adapter, enter an integer between 1 and 245.
CDB Transmit Queue Ring Size field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

Step 12 In the **Receive/Transmit Queues** area, review the information in the following fields:

Name	Description
FC Work Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 128.
FC Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 2048.

Step 13 In the **Boot Table** area, review the information in the following fields:

Name	Description
Index column	The unique identifier for the boot target.
Target WWPN column	The World Wide Port Name (WWPN) that corresponds to the location of the boot image.
LUN column	The LUN ID that corresponds to the location of the boot image.
Add Boot Entry button	Opens a dialog box that allows you to specify a new WWPN and LUN ID.
Edit Boot Entry button	Opens a dialog box that allows you to change the WWPN and LUN ID for the selected boot target.
Delete Boot Entry button	Deletes the selected boot target after you confirm the deletion.

Step 14 In the **Persistent Bindings** area, review the information in the following fields:

Name	Description
Index column	The unique identifier for the binding.
Target WWPN column	The target World Wide Port Name with which the binding is associated.
Host WWPN column	The host World Wide Port Name with which the binding is associated.
Bus ID column	The bus ID with which the binding is associated.
Target ID column	The target ID on the host system with which the binding is associated.
Rebuild Persistent Bindings button	Clears all unused bindings and resets the ones that are in use.

Modifying vHBA Properties

Step 1 In the **Navigation** pane, click the **Networking** menu.

Step 2 In the **Networking** pane, select the Adapter Card that you want to modify.

Step 3 In the **Adapter Card** pane, click the **vHBAs** tab.

Step 4 In the **vHBAs** pane, click **fc0** or **fc1**.

Step 5 In the **General** area, update the following fields:

Name	Description
Name field	The name of the virtual HBA. This name cannot be changed after the vHBA has been created.
Initiator WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
Initiator WWPNN field	The WWPNN associated with the vHBA. To let the system generate the WWPNN, select AUTO . To specify a WWPNN, click the second radio button and enter the WWPNN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
Persistent LUN Binding check box	If checked, any LUN ID associations are retained in memory until they are manually cleared.
Uplink Port drop-down list	The uplink port associated with the vHBA. Note This value cannot be changed for the system-defined vHBAs fc0 and fc1.
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
PCI Order field	The order in which this vHBA will be used. To let the system set the order, select ANY . To specify an order, select the second radio button and enter an integer between 0 and 17.

Name	Description
vHBA Type drop-down list	<p>Note This option is available only with 14xx series and VIC 15428 adapters.</p> <p>The vHBA type used in this policy. vHBAs supporting FC and FC-NVMe can now be created on the same adapter. The vHBA type used in this policy can be one of the following:</p> <ul style="list-style-type: none"> • fc-initiator—Legacy SCSI FC vHBA initiator • fc-target—vHBA that supports SCSI FC target functionality <p>Note This option is available as a Tech Preview.</p> <ul style="list-style-type: none"> • fc-nvme-initiator—vHBA that is an FC NVME initiator, which discovers FC NVME targets and connects to them. • fc-nvme-target—vHBA that acts as an FC NVME target and provides connectivity to the NVME storage.
Class of Service field	<p>The CoS for the vHBA.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Rate Limit field	<p>The data rate limit for traffic on this vHBA, in Mbps.</p> <p>If you want this vHBA to have an unlimited data rate, select OFF. Otherwise, click the second radio button and enter an integer between 1 and 10,000.</p> <p>Note This option cannot be used in VNTAG mode.</p>
EDTOV field	<p>The error detect timeout value (EDTOV), which is the number of milliseconds to wait before the system assumes that an error has occurred.</p> <p>Enter an integer between 1,000 and 100,000. The default is 2,000 milliseconds.</p>
RATOV field	<p>The resource allocation timeout value (RATOV), which is the number of milliseconds to wait before the system assumes that a resource cannot be properly allocated.</p> <p>Enter an integer between 5,000 and 100,000. The default is 10,000 milliseconds.</p>
Max Data Field Size field	<p>The maximum size of the Fibre Channel frame payload bytes that the vHBA supports.</p> <p>Enter an integer between 256 and 2112.</p>
Channel Number field	<p>The channel number that will be assigned to this vHBA.</p> <p>Enter an integer between 1 and 1,000.</p> <p>Note VNTAG mode is required for this option.</p>

Name	Description
PCI Link	It is read-only field.
Port Profile drop-down list	The port profile that should be associated with the vHBA, if any. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.

Step 6 In the **Error Recovery** area, update the following fields:

Name	Description
FCP Error Recovery check box	If checked, the system uses FCP Sequence Level Error Recovery protocol (FC-TAPE).
Link Down Timeout field	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240,000.
Port Down I/O Retry Count field	The number of times an I/O request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255.
IO Timeout Retry field	The time period till which the system waits for timeout before retrying. When a disk does not respond for I/O within the defined timeout period, the driver aborts the pending command, and resends the same I/O after the timer expires. Enter an integer between 1 and 59.
Port Down Timeout field	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. Enter an integer between 0 and 240,000.

Step 7 In the **Fibre Channel Interrupt** area, update the following fields:

Name	Description
Interrupt Mode drop-down list	The preferred driver interrupt mode. This can be one of the following: <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—MSI only. • INTx—PCI INTx interrupts.

Step 8 In the **Fibre Channel Port** area, update the following fields:

Name	Description
I/O Throttle Count field	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1,024.
LUNs per Target field	The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation. Enter an integer between 1 and 4,096.
LUN Queue Depth field	The number of commands that the HBA can send or receive in a single chunk per LUN. This parameter adjusts the initial queue depth for all LUNs on the adapter. Default value is 20 for physical miniports and 250 for virtual miniports.

Step 9 In the **Fibre Channel Port FLOGI** area, update the following fields:

Name	Description
FLOGI Retries field	The number of times that the system tries to log in to the fabric after the first failure. To specify an unlimited number of retries, select the INFINITE radio button. Otherwise select the second radio button and enter an integer into the corresponding field.
FLOGI Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1,000 and 255,000.

Step 10 In the **Fibre Channel Port PLOGI** area, update the following fields:

Name	Description
PLOGI Retries field	The number of times that the system tries to log in to a port after the first failure. Enter an integer between 0 and 255.
PLOGI Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1,000 and 255,000.

Step 11 In the **SCSI I/O** area, update the following fields:

Name	Description
CDB Transmit Queue Count field	The number of SCSI I/O queue resources the system should allocate. For Cisco UCS VIC 14xx series adapters, enter an integer between 1 and 64. For any other VIC adapter, enter an integer between 1 and 245.
CDB Transmit Queue Ring Size field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

Step 12 In the **Receive/Transmit Queues** area, update the following fields:

Name	Description
FC Work Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 128.
FC Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 2048.

Step 13 Click **Save Changes**.

Step 14 In the **Boot Table** area, update the following fields or add a new entry:

Name	Description
Index column	The unique identifier for the boot target.
Target WWPN column	The World Wide Port Name (WWPN) that corresponds to the location of the boot image.
LUN column	The LUN ID that corresponds to the location of the boot image.
Add Boot Entry button	Opens a dialog box that allows you to specify a new WWPN and LUN ID.
Edit Boot Entry button	Opens a dialog box that allows you to change the WWPN and LUN ID for the selected boot target.
Delete Boot Entry button	Deletes the selected boot target after you confirm the deletion.

Step 15 In the **Persistent Bindings** area, update the following fields or add a new entry:

Name	Description
Index column	The unique identifier for the binding.
Target WWPN column	The target World Wide Port Name with which the binding is associated.
Host WWPN column	The host World Wide Port Name with which the binding is associated.
Bus ID column	The bus ID with which the binding is associated.
Target ID column	The target ID on the host system with which the binding is associated.
Rebuild Persistent Bindings button	Clears all unused bindings and resets the ones that are in use.

Creating a vHBA

The Cisco UCS Virtual Interface Cards provide two vHBAs and two vNICs by default. You can create up to 14 additional vHBAs or vNICs on these adapter cards.

Cisco UCS 1455, 1457, and 1467 Virtual Interface Cards, in non-port channel mode, provide four vHBAs and four vNICs by default. You can create up to 10 additional vHBAs or vNICs on these adapter cards.

Before you begin

Ensure that **Enable VNTAG Mode** under **Adapter Card Properties** in the **General** tab is checked.

-
- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** In the **Host Fibre Channel Interfaces** area, choose one of these actions:
- To create a vHBA using default configuration settings, click **Add vHBA**.
 - To create a vHBA using the same configuration settings as an existing vHBA, select that vHBA and click **Clone vHBA**.
- The **Add vHBA** dialog box appears.
- Step 5** In the **Add vHBA** dialog box, enter a name for the vHBA in the **Name** entry box.
- Step 6** Configure the new vHBA as described in [Modifying vHBA Properties, on page 16](#).
- Step 7** Click **Add vHBA**.
-

What to do next

- Reboot the server to create the vHBA.

Deleting a vHBA

Default vHBAs cannot be deleted. You can delete any other vHBAs created using VNTAG mode.

-
- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** In the **Host Fibre Channel Interfaces** area, select a vHBA or vHBAs from the table.
- Note** You cannot delete either of the two default vHBAs, **fc0** or **fc1**.
- Step 5** Click **Delete vHBAs** and click **OK** to confirm.
-

What to do next

Reboot the server to delete the vHBA.

vHBA Boot Table

In the vHBA boot table, you can specify up to four LUNs from which the server can boot.

Creating a Boot Table Entry

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** Select a vHBA from the list of available vHBAs under **vHBAs** in the **vHBAs** tab.
The related **vHBA Properties** pane is displayed in the right hand side of the window.
- Step 5** Scroll down to view **Boot Table** at the end of the page.
- Step 6** Click the **Add Boot Entry** button to open the **Add Boot Entry** dialog box.
- Step 7** In the **Add Boot Entry** dialog box, review the following information and perform the actions specified:

Name	Description
Index field	The default value for this field is 0 .
Target WWPN field	The World Wide Port Name (WWPN) that corresponds to the location of the boot image. Enter the WWPN in the format hh : hh : hh : hh : hh : hh : hh : hh .
LUN ID field	The LUN ID that corresponds to the location of the boot image. Enter an ID between 0 and 255.
Add Boot Entry button	Adds the specified location to the boot table.
Reset Values button	Clears the values currently entered in the fields.
Cancel button	Closes the dialog box without saving any changes made while the dialog box was open.

Deleting a Boot Table Entry

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** Select a vHBA from the list of available vHBAs under **vHBAs** in the **vHBAs** tab.
The related **vHBA Properties** pane is displayed in the right hand side of the window.
- Step 5** Scroll down to view **Boot Table** at the end of the page.
- Step 6** In the **Boot Table** area, click the entry to be deleted.
- Step 7** Click **Delete Boot Entry** and click **OK** to confirm.

vHBA Persistent Binding

Persistent binding ensures that the system-assigned mapping of Fibre Channel targets is maintained after a reboot.

Viewing Persistent Bindings

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** In the **vHBAs** pane, click **fc0** or **fc1**.
- Step 5** In the **Persistent Bindings** area, review the following information:

Name	Description
Index column	The unique identifier for the binding.
Target WWPN column	The target World Wide Port Name with which the binding is associated.
Host WWPN column	The host World Wide Port Name with which the binding is associated.
Bus ID column	The bus ID with which the binding is associated.
Target ID column	The target ID on the host system with which the binding is associated.
Rebuild Persistent Bindings button	Clears all unused bindings and resets the ones that are in use.

Rebuilding Persistent Bindings

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 4** In the **vHBAs** pane, click **fc0** or **fc1**.
- Step 5** Scroll down to the **Persistent Bindings** area.
- Step 6** Click the **Rebuild Persistent Bindings** button.
- Step 7** Click **OK** to confirm.

Managing vNICs

Guidelines for Managing vNICs

When managing vNICs, consider the following guidelines and restrictions:

- The Cisco UCS Virtual Interface Cards provide two vHBAs and two vNICs by default. You can create up to 14 additional vHBAs or vNICs on these adapter cards.

Additional vHBAs can be created using VNTAG mode.

The Cisco UCS 1455, 1457, and 1467 Virtual Interface Cards, in non-port channel mode, provide four vHBAs and four vNICs by default. You can create up to 10 additional vHBAs or vNICs on these adapter cards.



Note If VNTAG mode is enabled for the adapter, you must assign a channel number to a vNIC when you create it.

- After making configuration changes, you must reboot the host for settings to take effect.

Viewing vNIC Properties

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to view.
- Step 3** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 4** In the **vNICs** pane, click **eth0** or **eth1**.
- Step 5** In the **General** area under **vNIC Properties** area, review the following fields:

General Area

Name	Description
Name field	The name for the virtual NIC. This name cannot be changed after the vNIC has been created.
CDN field	The Consistent Device Name (CDN) that you can assign to the ethernet vNICs on the VIC cards. Assigning a specific CDN to a device helps in identifying it on the host OS. Note This feature works only when the CDN Support for VIC token is enabled in the BIOS.
MTU field	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9000.

Name	Description
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	<p>The MAC address associated with the vNIC.</p> <p>To let the adapter select an available MAC address from its internal pool, select Auto. To specify an address, click the second radio button and enter the MAC address in the corresponding field.</p>
Class of Service field	<p>The class of service to associate with traffic from this vNIC.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Trust Host CoS check box	Check this box if you want the vNIC to use the class of service provided by the host operating system.
PCI Order field	<p>The order in which this vNIC will be used.</p> <p>To specify an order, enter an integer within the displayed range.</p>
Default VLAN radio button	<p>If there is no default VLAN for this vNIC, select NONE. Otherwise, select the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p>Note This option cannot be used in VNTAG mode.</p>
VLAN Mode drop-down list	<p>If you want to use VLAN trunking, select TRUNK. Otherwise, select ACCESS.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Enable PTP check box	<p>Check this box to enable Precision Time Protocol (PTP).</p> <p>Precision Time Protocol (PTP) precisely synchronizes the server clock with other devices and peripherals on Linux operating systems.</p> <p>Clocks managed by PTP follow a client-worker hierarchy, with workers synchronized to a master client. The hierarchy is updated by the best master clock (BMC) algorithm, which runs on every clock. One PTP interface per adapter must be enabled to synchronize it to the grand master clock.</p> <p>Note</p> <ul style="list-style-type: none"> • This option is supported only with Linux operating system. • This option is available only for Cisco UCS VIC 15xxx series adapters. <p style="padding-left: 40px;">This option is available only on some Cisco UCS C-Series servers.</p> <ul style="list-style-type: none"> • For the PTP enabling to be effective, it is required that you reboot the server.

Name	Description
Rate Limit radio button	<p>If you want this vNIC to have an unlimited data rate, select OFF. Otherwise, click the second radio button and enter a rate limit in the associated field.</p> <p>Enter an integer between 1 and 10,000 Mbps.</p> <p>For VIC 13xx controllers, you can enter an integer between 1 and 40,000 Mbps.</p> <p>For VIC 1455, 1457, and 1467 controllers:</p> <ul style="list-style-type: none"> • If the adapter is connected to 25 Gbps link on a Switch, then you can enter an integer between 1 to 25,000 Mbps for the Rate Limit field. • If the adapter is connected to 10 Gbps link on a Switch, then you can enter an integer between 1 to 10,000 Mbps for the Rate Limit field. <p>For VIC 1495, 1497, and 1477 controllers:</p> <ul style="list-style-type: none"> • If the adapter is connected to 40 Gbps link on a switch, then you can enter an integer between 1 to 40,000 Mbps for the Rate Limit field. • If the adapter is connected to 100 Gbps link on a switch, then you can enter an integer between 1 to 100,000 Mbps for the Rate Limit field. <p>Note This option cannot be used in VNTAG mode.</p>
Channel Number field	<p>Select the channel number that will be assigned to this vNIC.</p> <p>Note VNTAG mode is required for this option.</p>
PCI Link field	<p>The link through which vNICs can be connected. These are the following values:</p> <ul style="list-style-type: none"> • 0 - The first cross-edged link where the vNIC is placed. • 1 - The second cross-edged link where the vNIC is placed. <p>Note</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers.
Enable NVGRE check box	<p>Check this box to enable Network Virtualization using Generic Routing Encapsulation.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385 cards.
Enable VXLAN check box	<p>Check this box to enable Virtual Extensible LAN.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385, VIC 14xx, and VIC 15xxx.

Name	Description
Geneve Offload check box	<p>Beginning with release 4.1(2a), Cisco IMC supports Generic Network Virtualization Encapsulation (Geneve) Offload feature with Cisco VIC 14xx series and VIC 15xxx adapters in ESX 7.0 (NSX-T 3.0) and ESX 6.7U3(NSX-T 2.5) OS.</p> <p>Geneve is a tunnel encapsulation functionality for network traffic . Check this box if you want to enable Geneve Offload encapsulation in Cisco VIC 14xx series adapters.</p> <p>Un-check this box to disable Geneve Offload, in order to prevent non-encapsulated UDP packets whose destination port numbers match with the Geneve destination port from being treated as tunneled packets.</p> <p>If you enable Geneve Offload feature, then Cisco recommends the following settings:</p> <ul style="list-style-type: none"> • Transmit Queue Count—1 • Transmit Queue Ring Size—4096 • Receive Queue Count—8 • Receive Queue Ring Size—4096 • Completion Queue Count—9 • Interrupt Count—11 <p>Note You cannot enable the following when Geneve Offload is enabled in a setup with Cisco VIC 14xx series:</p> <ul style="list-style-type: none"> • RDMA on the same vNIC • usNIC on the same vNIC • Non-Port Channel Mode on Cisco VIC 145x adapters • aRFS • Advanced Filters • NetQueue <p>Note You cannot enable the following when Geneve Offload is enabled in a setup with Cisco VIC 15xxx:</p> <ul style="list-style-type: none"> • aRFS • RoCEv2 <p>Outer IPV6 is not supported with GENEVE Offload feature.</p> <p>Downgrade Limitation—If Geneve Offload is enabled, you cannot downgrade to any release earlier than 4.1(2a).</p>
Advanced Filter check box	Check this box to enable advanced filter options in vNICs.

Name	Description
Port Profile drop-down list	Select the port profile that should be associated with the vNIC. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Enable VMQ check box	Check this box to enable Virtual Machine Queue (VMQ).
Enable Multi Queue check box	Check this box to enable the Multi Queue option on vNICs. When enabled multi queue vNICs will be available to the host. By default this is disabled. Note <ul style="list-style-type: none"> • Multi queue is supported only on C-Series servers with 14xx and VIC 15xxx adapters. • VMQ must be in enabled state to enable this option. • When you enable this option on one of the vNICs, configuring only VMQ (without choosing multi-queue) on other vNICs is not supported. • When this option is enabled usNIC configuration will be disabled.
No. of Sub vNICs field	Number of sub vNICs available to the host when the multi queue option is enabled.
Enable aRFS check box	Check this box to enable Accelerated Receive Flow steering (aRFS). This option is available only on some Cisco UCS C-Series servers.
Enable Uplink Failover check box	Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems. Note VNTAG mode is required for this option.
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. Note VNTAG mode is required for this option.

Step 6

In the **Ethernet Interrupt** area, review the information in the following fields:

Name	Description
Interrupt Count 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 1024.

Name	Description
Interrupt Mode drop-down list	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.
Coalescing Time field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the Coalescing Time field.

Step 7

In the **TCP Offload** area, review the information in the following fields:

Name	Description
Enable Large Receive check box	If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput. If cleared, the CPU processes all large packets.
Enable TCP Rx Offload Checksum Validation check box	If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. If cleared, the CPU validates all packet checksums.
Enable TCP Segmentation Offload check box	If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. If cleared, the CPU segments large packets. Note This option is also known as Large Send Offload (LSO).
Enable TCP Tx Offload Checksum Generation check box	If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. If cleared, the CPU calculates all packet checksums.

Step 8

In the **Receive Side Scaling** area, review the information in the following fields:

Name	Description
Enable TCP Receive Side Scaling check box	Receive Side Scaling (RSS) distributes network receive processing across multiple CPUs in multiprocessor systems. If checked, network receive processing is shared across processors whenever possible. If cleared, network receive processing is always handled by a single processor even if additional processors are available.
Enable IPv4 RSS check box	If checked, RSS is enabled on IPv4 networks.
Enable TCP-IPv4 RSS check box	If checked, RSS is enabled for TCP transmissions across IPv4 networks.
Enable IPv6 RSS check box	If checked, RSS is enabled on IPv6 networks.
Enable TCP-IPv6 RSS check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.
Enable IPv6 Extension RSS check box	If checked, RSS is enabled for IPv6 extensions.
Enable TCP-IPv6 Extension RSS check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.

Step 9

In the **Ethernet Receive Queue** area, review the information in the following fields:

Name	Description
Count 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.

Step 10

In the **Ethernet Transmit Queue** area, review the information in the following fields:

Name	Description
Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.

Step 11

In the **Completion Queue** area, review the information in the following fields:

Name	Description
Count 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.

Name	Description
Ring Size	The number of descriptors in each completion queue. This value cannot be changed.

Step 12 In the **Multi Queue** area, review the information in the following fields:

Name	Description
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 1000.
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 1000.
Completion Queue Count field	The number of completion queue resources to allocate. Enter an integer between 1 and 2000.

Step 13 In the **RoCE Properties** area, review the information in the following fields:

Name	Description
RoCE checkbox	Check the check box to change the RoCE Properties.
Queue Pairs field	The number of queue pairs per adapter. Enter an integer between 1 and 2048. We recommend that this number be an integer power of 2. The recommended value for queue pairs per vNIC is 2048.
Memory Regions field	The number of memory regions per adapter. Enter an integer between 1 and 524288. We recommend that this number be an integer power of 2. The recommended value is 131072. The number of memory regions supported should be enough to meet application requirements as the regions are primarily used to send operation channel semantics.
Resource Groups field	The number of resource groups per adapter. Enter an integer between 1 and 128. We recommend 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. The resource group defines the total number of hardware resources such as WQ, RQ, CQ, and interrupts required to support the RDMA functionality, and is based on the total number of processor cores available with the host. The host chooses to dedicate a particular resource group to a core to maximize performance and get a better non-uniform memory access.

Name	Description
Class of Service drop-down list	NO Drop QOS COS to be specified. This same value should be configured at the up link switch. Default No Drop QOS COS is 5. Note This option is available only on some adapters.

What to do next

Reboot the server to create the vHBA.

Modifying vNIC Properties

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 4** In the **vNICs** pane, click **eth0** or **eth1**.
- Step 5** In the **General** area under **vNIC Properties** in the **vNICs** pane, update the following fields:

Name	Description
Name field	The name for the virtual NIC. This name cannot be changed after the vNIC has been created.
CDN field	The Consistent Device Name (CDN) that you can assign to the ethernet vNICs on the VIC cards. Assigning a specific CDN to a device helps in identifying it on the host OS. Note This feature works only when the CDN Support for VIC token is enabled in the BIOS.
MTU field	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9000.
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	The MAC address associated with the vNIC. To let the adapter select an available MAC address from its internal pool, select Auto . To specify an address, click the second radio button and enter the MAC address in the corresponding field.

Name	Description
Class of Service field	<p>The class of service to associate with traffic from this vNIC.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Trust Host CoS check box	<p>Check this box if you want the vNIC to use the class of service provided by the host operating system.</p>
PCI Order field	<p>The order in which this vNIC will be used.</p> <p>To specify an order, enter an integer within the displayed range.</p>
Default VLAN radio button	<p>If there is no default VLAN for this vNIC, select NONE. Otherwise, select the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p>Note This option cannot be used in VNTAG mode.</p>
VLAN Mode drop-down list	<p>If you want to use VLAN trunking, select TRUNK. Otherwise, select ACCESS.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Enable PTP check box	<p>Check this box to enable Precision Time Protocol (PTP).</p> <p>Precision Time Protocol (PTP) precisely synchronizes the server clock with other devices and peripherals on Linux operating systems.</p> <p>Clocks managed by PTP follow a client-worker hierarchy, with workers synchronized to a master client. The hierarchy is updated by the best master clock (BMC) algorithm, which runs on every clock. One PTP interface per adapter must be enabled to synchronize it to the grand master clock.</p> <p>Note</p> <ul style="list-style-type: none"> • This option is supported only with Linux operating system. • This option is available only for Cisco UCS VIC 15xxx series adapters. <p>This option is available only on some Cisco UCS C-Series servers.</p> <ul style="list-style-type: none"> • For the PTP enabling to be effective, it is required that you reboot the server.

Name	Description
Rate Limit radio button	<p>If you want this vNIC to have an unlimited data rate, select OFF. Otherwise, click the second radio button and enter a rate limit in the associated field.</p> <p>Enter an integer between 1 and 10,000 Mbps.</p> <p>For VIC 13xx controllers, you can enter an integer between 1 and 40,000 Mbps.</p> <p>For VIC 1455, 1457, and 1467 controllers:</p> <ul style="list-style-type: none"> • If the adapter is connected to 25 Gbps link on a Switch, then you can enter an integer between 1 to 25,000 Mbps for the Rate Limit field. • If the adapter is connected to 10 Gbps link on a Switch, then you can enter an integer between 1 to 10,000 Mbps for the Rate Limit field. <p>For VIC 1495, 1497, and 1477 controllers:</p> <ul style="list-style-type: none"> • If the adapter is connected to 40 Gbps link on a switch, then you can enter an integer between 1 to 40,000 Mbps for the Rate Limit field. • If the adapter is connected to 100 Gbps link on a switch, then you can enter an integer between 1 to 100,000 Mbps for the Rate Limit field. <p>Note This option cannot be used in VNTAG mode.</p>
Channel Number field	<p>Select the channel number that will be assigned to this vNIC.</p> <p>Note VNTAG mode is required for this option.</p>
PCI Link field	<p>The link through which vNICs can be connected. These are the following values:</p> <ul style="list-style-type: none"> • 0 - The first cross-edged link where the vNIC is placed. • 1 - The second cross-edged link where the vNIC is placed. <p>Note</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers.
Enable NVGRE check box	<p>Check this box to enable Network Virtualization using Generic Routing Encapsulation.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385 cards.
Enable VXLAN check box	<p>Check this box to enable Virtual Extensible LAN.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385, VIC 14xx, and VIC 15xxx.

Name	Description
Geneve Offload check box	<p>Beginning with release 4.1(2a), Cisco IMC supports Generic Network Virtualization Encapsulation (Geneve) Offload feature with Cisco VIC 14xx series and VIC 15xxx adapters in ESX 7.0 (NSX-T 3.0) and ESX 6.7U3(NSX-T 2.5) OS.</p> <p>Geneve is a tunnel encapsulation functionality for network traffic . Check this box if you want to enable Geneve Offload encapsulation in Cisco VIC 14xx series adapters.</p> <p>Un-check this box to disable Geneve Offload, in order to prevent non-encapsulated UDP packets whose destination port numbers match with the Geneve destination port from being treated as tunneled packets.</p> <p>If you enable Geneve Offload feature, then Cisco recommends the following settings:</p> <ul style="list-style-type: none"> • Transmit Queue Count—1 • Transmit Queue Ring Size—4096 • Receive Queue Count—8 • Receive Queue Ring Size—4096 • Completion Queue Count—9 • Interrupt Count—11 <p>Note You cannot enable the following when Geneve Offload is enabled in a setup with Cisco VIC 14xx series:</p> <ul style="list-style-type: none"> • RDMA on the same vNIC • usNIC on the same vNIC • Non-Port Channel Mode on Cisco VIC 145x adapters • aRFS • Advanced Filters • NetQueue <p>Note You cannot enable the following when Geneve Offload is enabled in a setup with Cisco VIC 15xxx:</p> <ul style="list-style-type: none"> • aRFS • RoCEv2 <p>Outer IPV6 is not supported with GENEVE Offload feature.</p> <p>Downgrade Limitation—If Geneve Offload is enabled, you cannot downgrade to any release earlier than 4.1(2a).</p>
Advanced Filter check box	Check this box to enable advanced filter options in vNICs.

Name	Description
Port Profile drop-down list	Select the port profile that should be associated with the vNIC. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Enable VMQ check box	Check this box to enable Virtual Machine Queue (VMQ).
Enable Multi Queue check box	Check this box to enable the Multi Queue option on vNICs. When enabled multi queue vNICs will be available to the host. By default this is disabled. Note <ul style="list-style-type: none"> • Multi queue is supported only on C-Series servers with 14xx and VIC 15xxx adapters. • VMQ must be in enabled state to enable this option. • When you enable this option on one of the vNICs, configuring only VMQ (without choosing multi-queue) on other vNICs is not supported. • When this option is enabled usNIC configuration will be disabled.
No. of Sub vNICs field	Number of sub vNICs available to the host when the multi queue option is enabled.
Enable aRFS check box	Check this box to enable Accelerated Receive Flow steering (aRFS). This option is available only on some Cisco UCS C-Series servers.
Enable Uplink Failover check box	Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems. Note VNTAG mode is required for this option.
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. Note VNTAG mode is required for this option.

Step 6

In the **Ethernet Interrupt** area, update the following fields:

Name	Description
Interrupt Count 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 1024.

Name	Description
Coalescing Time field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time 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 Coalescing Time field.
Interrupt Mode drop-down list	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.

Step 7

In the **TCP Offload** area, update the following fields:

Name	Description
Enable Large Receive check box	If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput. If cleared, the CPU processes all large packets.
Enable TCP Segmentation Offload check box	If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. If cleared, the CPU segments large packets. Note This option is also known as Large Send Offload (LSO).
Enable TCP Rx Offload Checksum Validation check box	If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. If cleared, the CPU validates all packet checksums.
Enable TCP Tx Offload Checksum Generation check box	If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. If cleared, the CPU calculates all packet checksums.

Step 8

In the **Receive Side Scaling** area, update the following fields:

Name	Description
Enable TCP Receive Side Scaling check box	Receive Side Scaling (RSS) distributes network receive processing across multiple CPUs in multiprocessor systems. If checked, network receive processing is shared across processors whenever possible. If cleared, network receive processing is always handled by a single processor even if additional processors are available.
Enable IPv4 RSS check box	If checked, RSS is enabled on IPv4 networks.
Enable TCP-IPv4 RSS check box	If checked, RSS is enabled for TCP transmissions across IPv4 networks.
Enable IPv6 RSS check box	If checked, RSS is enabled on IPv6 networks.
Enable TCP-IPv6 RSS check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.
Enable IPv6 Extension RSS check box	If checked, RSS is enabled for IPv6 extensions.
Enable TCP-IPv6 Extension RSS check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.

Step 9

In the **Ethernet Receive Queue** area, update the following fields:

Name	Description
Count 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 16384. VIC 14xx Series adapters support a 4K (4096) maximum Ring Size VIC15xxx Series adapters support up to 16K Ring Size.

Step 10

In the **Ethernet Transmit Queue** area, update the following fields:

Name	Description
Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 16384. VIC 14xx Series adapters support a 4K (4096) maximum Ring Size VIC15xxx Series adapters support up to 16K Ring Size.

Step 11

In the **Completion Queue** area, update the following fields:

Name	Description
Count 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.
Ring Size	The number of descriptors in each completion queue. This value cannot be changed.

Step 12 In the **Multi Queue** area, update the following details:

Name	Description
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 1000.
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 1000.
Completion Queue Count 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 2000.

Step 13 In the **RoCE Properties** area, update the following fields:

Name	Description
RoCE check box	Check the check box to change the RoCE Properties. Note If Multi Queue RoCE is enabled, ensure that VMQ RoCE is also enabled.
Queue Pairs field	The number of queue pairs per adapter. Enter an integer between 1 and 2048. We recommend that this number be an integer power of 2.
Memory Regions field	The number of memory regions per adapter. Enter an integer between 1 and 524288. We recommend that this number be an integer power of 2.
Resource Groups field	The number of resource groups per adapter. Enter an integer between 1 and 128. We recommend 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.
Class of Service field	This field is read-only and is set to 5. Note This option is available only on some of the adapters.

Step 14 Click **Save Changes**.

What to do next

Reboot the server to modify the vNIC.

Creating a vNIC

The Cisco UCS Virtual Interface Cards provide two vHBAs and two vNICs by default. You can create up to 14 additional vHBAs or vNICs on these adapter cards.

The Cisco UCS 1455, 1457, and 1467 Virtual Interface Cards, in non-port channel mode, provide four vHBAs and four vNICs by default. You can create up to 10 additional vHBAs or vNICs on these adapter cards.

Step 1 In the **Navigation** pane, click the **Networking** menu.

Step 2 In the **Networking** pane, select the Adapter Card that you want to modify.

Step 3 In the **Adapter Card** pane, click the **vNICs** tab.

Step 4 In the **Host Ethernet Interfaces** area, choose one of these actions:

- To create a vNIC using default configuration settings, click **Add vNIC**.
- To create a vNIC using the same configuration settings as an existing vNIC, select that vNIC and click **Clone vNIC**.

The **Add vNIC** dialog box appears.

Step 5 In the **Add vNIC** dialog box, enter a name for the vNIC in the **Name** entry box.

Step 6 In the **Add vNIC** dialog box, enter a channel number for the vNIC in the **Channel Number** entry box.

Note If VNTAG is enabled on the adapter, you must assign a channel number for the vNIC when you create it.

Step 7 Click **Add vNIC**.

What to do next

If configuration changes are required, configure the new vNIC as described in [Modifying vNIC Properties, on page 33](#).

Deleting a vNIC

Step 1 In the **Navigation** pane, click the **Networking** menu.

Step 2 In the **Networking** pane, select the Adapter Card that you want to modify.

Step 3 In the **Adapter Card** pane, click the **vNICs** tab.

Step 4 In the **Host Ethernet Interfaces** area, select a vNIC from the table.

Note You cannot delete either of the two default vNICs, **eth0** or **eth1**.

Step 5 Click **Delete vNIC** and click **OK** to confirm.

Configuring iSCSI Boot Capability

Configuring iSCSI Boot Capability for vNICs

To configure the iSCSI boot capability on a vNIC:

- You must log in with admin privileges to perform this task.
- To configure a vNIC to boot a server remotely from an iSCSI storage target, you must enable the PXE boot option on the vNIC.



Note You can configure a maximum of 2 iSCSI vNICs for each host.

Configuring iSCSI Boot Capability on a vNIC

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 4** In the **vNICs** pane, click **eth0** or **eth1**.
- Step 5** Select the **iSCSI Boot Properties** area.
- Step 6** In the **General** area, update the following fields:

Name	Description
Name field	The name of the vNIC.
DHCP Network check box	Whether DHCP Network is enabled for the vNIC. If enabled, the initiator network configuration is obtained from the DHCP server.
DHCP iSCSI check box	Whether DHCP iSCSI is enabled for the vNIC. If enabled and the DHCP ID is set, the initiator IQN and target information are obtained from the DHCP server. Note If DHCP iSCSI is enabled without a DHCP ID, only the target information is obtained.
DHCP ID field	The vendor identifier string used by the adapter to obtain the initiator IQN and target information from the DHCP server. Enter a string up to 64 characters.
DHCP Timeout field	The number of seconds to wait before the initiator assumes that the DHCP server is unavailable. Enter an integer between 60 and 300 (default: 60 seconds)

Name	Description
Link Timeout field	The number of seconds to wait before the initiator assumes that the link is unavailable. Enter an integer between 0 and 255 (default: 15 seconds)
LUN Busy Retry Count field	The number of times to retry the connection in case of a failure during iSCSI LUN discovery. Enter an integer between 0 and 255. The default is 15.
IP Version field	The IP version to use during iSCSI boot.

Step 7

In the **Initiator** area, update the following fields:

Name	Description
Name field	A regular expression that defines the name of the iSCSI initiator. You can enter any alphanumeric string as well as the following special characters: <ul style="list-style-type: none"> • . (period) • : (colon) • - (dash) <p>Note The name is in the IQN format.</p>
IP Address field	The IP address of the iSCSI initiator.
Subnet Mask field	The subnet mask for the iSCSI initiator.
Gateway field	The default gateway.
Primary DNS field	The primary DNS server address.
Initiator Priority drop-down list	The initiator priority drop-down list.
Secondary DNS field	The secondary DNS server address.
TCP Timeout field	The number of seconds to wait before the initiator assumes that TCP is unavailable. Enter an integer between 0 and 255 (default: 15 seconds)
CHAP Name field	The Challenge-Handshake Authentication Protocol (CHAP) name of the initiator.
CHAP Secret field	The Challenge-Handshake Authentication Protocol (CHAP) shared secret of the initiator.

Step 8

In the **Primary Target** area, update the following fields:

Name	Description
Name field	The name of the primary target in the IQN format.
IP Address field	The IP address of the target.
TCP Port field	The TCP port associated with the target.
Boot LUN field	The Boot LUN associated with the target.
CHAP Name field	The Challenge-Handshake Authentication Protocol (CHAP) name of the initiator.
CHAP Secret field	The Challenge-Handshake Authentication Protocol (CHAP) shared secret of the initiator.

Step 9 In the **Secondary Target** area, update the following fields:

Name	Description
Name field	The name of the secondary target in the IQN format.
IP Address field	The IP address of the target.
TCP Port field	The TCP port associated with the target.
Boot LUN field	The Boot LUN associated with the target.
CHAP Name field	The Challenge-Handshake Authentication Protocol (CHAP) name of the initiator.
CHAP Secret field	The Challenge-Handshake Authentication Protocol (CHAP) shared secret of the initiator.

Name	Description
Configure iSCSI button	Configures iSCSI boot on the selected vNIC.
Unconfigure iSCSI button	Removes the configuration from the selected vNIC.
Reset Values button	Restores the values for the vNIC to the settings that were in effect when this dialog box was first opened.
Cancel button	Closes the dialog box without making any changes.

Step 10 Click **Save Changes**.

Removing iSCSI Boot Configuration from a vNIC

Before you begin

You must log in with admin privileges to perform this task.

-
- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 4** In the **vNICs** pane, click **eth0** or **eth1**.
- Step 5** Select the **iSCSI Boot Properties** area.
- Step 6** Click the **Unconfigure iSCSI Boot** button at the bottom of the **iSCSI Boot Properties** area.
-

What to do next

Reboot the server to remove the iSCSI Boot Configuration.

Managing Cisco usNIC

Overview of Cisco usNIC

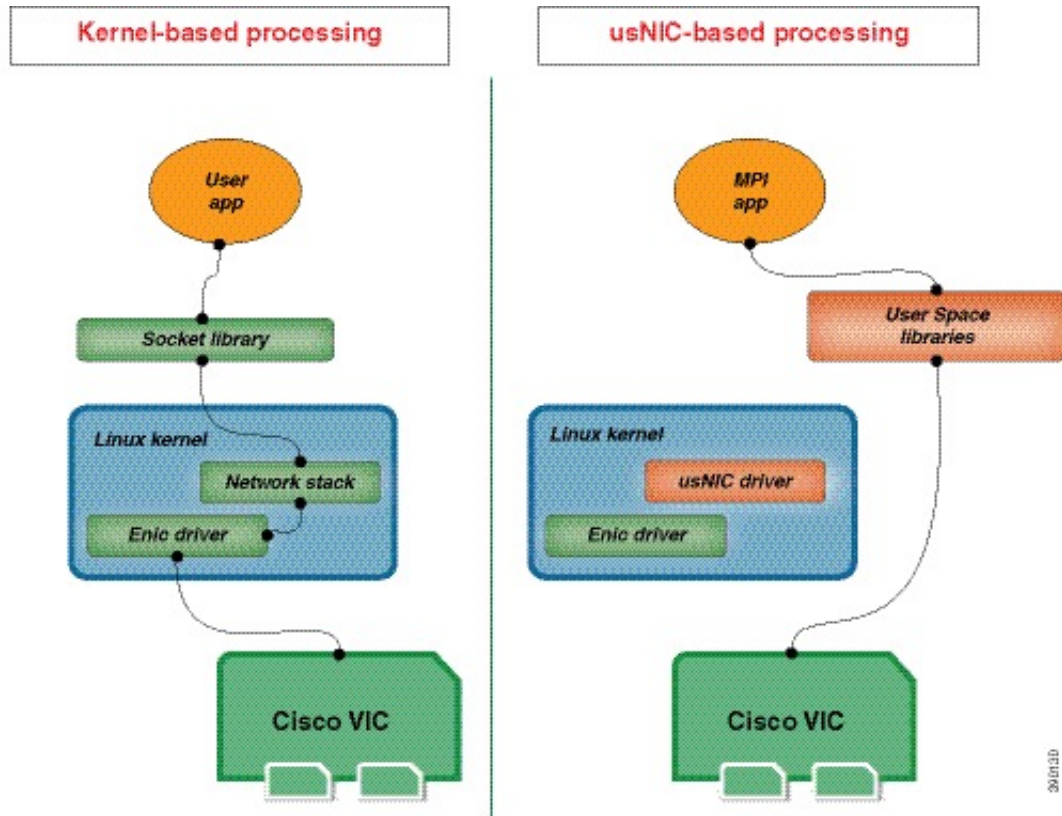
The Cisco user-space NIC (Cisco usNIC) feature improves the performance of software applications that run on the Cisco UCS servers in your data center by bypassing the kernel when sending and receiving networking packets. The applications interact directly with a Cisco UCS VIC second generation or later generation adapter, such as the , which improves the networking performance of your high-performance computing cluster. To benefit from Cisco usNIC, your applications must use the Message Passing Interface (MPI) instead of sockets or other communication APIs.

Cisco usNIC offers the following benefits for your MPI applications:

- Provides a low-latency and high-throughput communication transport.
- Employs the standard and application-independent Ethernet protocol.
- Takes advantage of lowlatency forwarding, Unified Fabric, and integrated management support in the following Cisco data center platforms:
 - Cisco UCS server
 - Cisco UCS VIC second generation or later generation adapter
 - 10 or 40GbE networks

Standard Ethernet applications use user-space socket libraries, which invoke the networking stack in the Linux kernel. The networking stack then uses the Cisco eNIC driver to communicate with the Cisco VIC hardware. The following figure shows the contrast between a regular software application and an MPI application that uses Cisco usNIC.

Figure 1: Kernel-Based Network Communication versus Cisco usNIC-Based Communication



Viewing and Configuring Cisco usNIC using the Cisco IMC GUI

Before you begin

You must log in to the Cisco IMC GUI with admin privileges to perform this task. Click Play on this [video](#) to watch how to configure Cisco usNIC in Cisco IMC.

- Step 1** Log into the Cisco IMC GUI.
For more information about how to log into Cisco IMC, see [Cisco UCS C-Series Servers Integrated Management Controller GUI Configuration Guide](#).
- Step 2** In the **Navigation** pane, click the **Networking** menu.
- Step 3** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 4** In the **Adapter Card** pane, click the vNICs tab.
- Step 5** In the vNICs pane, click **eth0** or **eth1**.
- Step 6** In the **usNIC** area, review and update the following fields.

Name	Description
Name	The name for the vNIC that is the parent of the usNIC.
	Note This field is read-only.

Name	Description
usNIC field	<p>The number of usNICs assigned to the specific vNIC.</p> <p>Enter an integer between 0 and 225.</p> <p>To assign additional usNICs to a specified vNIC, enter value higher than the existing value.</p> <p>To delete usNICs from a specified vNIC, enter value smaller than the existing value.</p> <p>To delete all the usNICs assigned to a vNIC, enter zero.</p>
Transmit Queue Count field	<p>The number of transmit queue resources to allocate.</p> <p>Enter an integer between 1 and 256.</p>
Receive Queue Count field	<p>The number of receive queue resources to allocate.</p> <p>Enter an integer between 1 and 256.</p>
Completion Queue Count field	<p>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.</p> <p>Enter an integer between 1 and 512.</p>
Transmit Queue Ring Size field	<p>The number of descriptors in each transmit queue.</p> <p>Enter an integer between 64 and 4096.</p>
Receive Queue Ring Size field	<p>The number of descriptors in each receive queue.</p> <p>Enter an integer between 64 and 4096.</p>
Interrupt Count field	<p>The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources.</p> <p>Enter an integer between 1 and 514.</p>
Interrupt Coalescing Type drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the Coalescing Time field.
Interrupt Coalescing Timer Time field	<p>The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent.</p> <p>Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.</p>

Name	Description
Class of Service field	<p>The class of service to associate with traffic from this usNIC.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
TCP Segment Offload check box	<p>If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate.</p> <p>If cleared, the CPU segments large packets.</p> <p>Note This option is also known as Large Send Offload (LSO).</p>
Large Receive check box	<p>If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput.</p> <p>If cleared, the CPU processes all large packets.</p>
TCP Tx Checksum check box	<p>If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead.</p> <p>If cleared, the CPU calculates all packet checksums.</p>
TCP Rx Checksum check box	<p>If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead.</p> <p>If cleared, the CPU validates all packet checksums.</p>

Step 7 Click **Save Changes**.

The changes take effect upon the next server reboot.

Viewing usNIC Properties

Step 1 In the **Navigation** pane, click the **Networking** menu.

Step 2 In the **Networking** pane, select the Adapter Card that you want to view.

Step 3 In the **Adapter Card** pane, click the **vNICs** tab.

Step 4 In the **vNICs** pane, click **eth0** or **eth1**.

Step 5 In the **Host Ethernet Interfaces** pane's **usNIC Properties** area, review the information in the following fields:

Name	Description
Name	The name for the vNIC that is the parent of the usNIC. Note This field is read-only.
usNIC field	The number of usNICs assigned to the specific vNIC. Enter an integer between 0 and 225. To assign additional usNICs to a specified vNIC, enter value higher than the existing value. To delete usNICs from a specified vNIC, enter value smaller than the existing value. To delete all the usNICs assigned to a vNIC, enter zero.
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
Completion Queue Count 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.
Transmit Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.
Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.
Interrupt Count 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.
Interrupt Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time 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 Coalescing Time field.

Name	Description
Interrupt Coalescing Timer Time field	<p>The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent.</p> <p>Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.</p>
Class of Service field	<p>The class of service to associate with traffic from this usNIC.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>
TCP Segment Offload check box	<p>If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate.</p> <p>If cleared, the CPU segments large packets.</p> <p>Note This option is also known as Large Send Offload (LSO).</p>
Large Receive check box	<p>If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput.</p> <p>If cleared, the CPU processes all large packets.</p>
TCP Tx Checksum check box	<p>If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead.</p> <p>If cleared, the CPU calculates all packet checksums.</p>
TCP Rx Checksum check box	<p>If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead.</p> <p>If cleared, the CPU validates all packet checksums.</p>

Backing Up and Restoring the Adapter Configuration

Exporting the Adapter Configuration

The adapter configuration can be exported as an XML file to a remote server which can be one of the following:

- TFTP

- FTP
- SFTP
- SCP
- HTTP

Before you begin

Obtain the remote server IP address.

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** In the **Adapter Card** pane, click the **General** tab.
- Step 4** In the **Actions** area of the **General** tab, click **Export vNIC**.

The **Export vNIC** dialog box opens.

- Step 5** In the **Export Adapter Configuration** dialog box, update the following fields:

Name	Description
Export To drop-down list	<p>The remote server type. This can be one of the following:</p> <ul style="list-style-type: none"> • TFTP Server • FTP Server • SFTP Server • SCP Server • HTTP Server <p>Note If you chose SCP or SFTP as the remote server type while performing this action, a pop-up window is displayed with the message <i>Server (RSA) key fingerprint is <server_finger_print_ID> Do you wish to continue?</i>. Click Yes or No depending on the authenticity of the server fingerprint.</p> <p>The fingerprint is based on the host's public key and helps you to identify or verify the host you are connecting to.</p>
IP Address or Host Name field	<p>The IPv4 or IPv6 address, or hostname of the server to which the adapter configuration file will be exported. Depending on the setting in the Export to drop-down list, the name of the field may vary.</p>
Path and Filename field	<p>The path and filename Cisco IMC should use when exporting the file to the remote server.</p>
Username	<p>The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or HTTP.</p>

Name	Description
Password	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.

Step 6 Click **Export vNIC**.

Importing the Adapter Configuration

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to modify.
- Step 3** Select the **General** tab.
- Step 4** In the **Actions** area of the **General** tab, click **Import vNIC**.
The **Import vNIC** dialog box is displayed.
- Step 5** In the **Import vNIC** dialog box, update the following fields:

Name	Description
Import from drop-down list	<p>The remote server type. This can be one of the following:</p> <ul style="list-style-type: none"> • TFTP Server • FTP Server • SFTP Server • SCP Server • HTTP Server <p>Note If you chose SCP or SFTP as the remote server type while performing this action, a pop-up window is displayed with the message <i>Server (RSA) key fingerprint is <server_finger_print_ID> Do you wish to continue?</i>. Click Yes or No depending on the authenticity of the server fingerprint.</p> <p>The fingerprint is based on the host's public key and helps you to identify or verify the host you are connecting to.</p>
IP Address or Host Name field	The IPv4 or IPv6 address, or hostname of the server on which the adapter configuration file resides. Depending on the setting in the Import from drop-down list, the name of the field may vary.
Path and Filename field	The path and filename of the configuration file on the remote server.
Username	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or HTTP.

Name	Description
Password	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.

Step 6 Click **Import vNIC**.

What to do next

Reboot the server to apply the imported configuration.

Restoring Adapter Defaults

- Step 1** In the **Navigation** pane, click the **Networking** menu.
 - Step 2** In the **Networking** pane, select the Adapter Card that you want to restore to default settings.
 - Step 3** Select the **General** tab.
 - Step 4** In the **Actions** area of the **General** tab, click **Reset To Defaults** and click **OK** to confirm.
-

Resetting the Adapter

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** pane, select the Adapter Card that you want to reset.
- Step 3** Select the **General** tab.
- Step 4** In the **Actions** area of the **General** tab, click **Reset** and click **OK** to confirm.

Note Resetting the adapter also resets the host.
