



# Managing Network Adapters

---

This chapter includes the following sections:

- [Overview of the Cisco UCS C-Series Network Adapters, page 1](#)
- [Viewing Network Adapter Properties, page 4](#)
- [Viewing VIC Adapter Properties, page 5](#)
- [Viewing Storage Adapter Properties, page 10](#)
- [Managing vHBAs, page 11](#)
- [Managing vNICs, page 24](#)
- [Backing Up and Restoring the Adapter Configuration, page 49](#)
- [Managing Adapter Firmware, page 52](#)
- [Resetting the Adapter, page 55](#)

## 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 P81E Virtual Interface Card
- Cisco UCS VIC 1225 Virtual Interface Card
- Cisco UCS VIC 1385 Virtual Interface Card
- Cisco UCS VIC 1227T Virtual Interface Card
- Cisco UCS VIC 1387 Virtual Interface Card

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 P81E Virtual Interface Card

The Cisco UCS P81E Virtual Interface Card is optimized for virtualized environments, for organizations that seek increased mobility in their physical environments, and for data centers that want reduced costs through NIC, HBA, cabling, and switch reduction and reduced management overhead. This Fibre Channel over Ethernet (FCoE) PCIe card offers the following benefits:

- Allows up to 16 virtual Fibre Channel and 16 virtual Ethernet adapters to be provisioned in virtualized or nonvirtualized environments using just-in-time provisioning, providing tremendous system flexibility and allowing consolidation of multiple physical adapters.
- Delivers uncompromising virtualization support, including hardware-based implementation of Cisco VN-Link technology and pass-through switching.
- Improves system security and manageability by providing visibility and portability of network policies and security all the way to the virtual machine.

The virtual interface card makes Cisco VN-Link connections to the parent fabric interconnects, which allows virtual links to connect virtual NICs in virtual machines to virtual interfaces in the interconnect. In a Cisco Unified Computing System environment, virtual links then can be managed, network profiles applied, and interfaces dynamically reprovisioned as virtual machines move between servers in the system.

### 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.

The Cisco UCS VIC 1225 implements the Cisco Virtual Machine Fabric Extender (VM-FEX), which unifies virtual and physical networking into a single infrastructure. It provides virtual-machine visibility from the physical network and a consistent network operations model for physical and virtual servers. In virtualized environments, this highly configurable and self-virtualized adapter provides integrated, modular LAN interfaces on Cisco UCS C-Series Rack-Mount Servers. Additional features and capabilities include:

- Supports up to 256 PCIe virtual devices, either virtual network interface cards (vNICs) or virtual host bus adapters (vHBAs), with high I/O operations per second (IOPS), support for lossless Ethernet, and 20 Gbps to servers.
- PCIe Gen2 x16 helps assure optimal bandwidth to the host for network-intensive applications with a redundant path to the fabric interconnect.
- Half-height design reserves full-height slots in servers for Cisco certified third-party adapters.
- Centrally managed by Cisco UCS Manager with support for Microsoft Windows, Red Hat Enterprise Linux, SUSE Linux, VMware vSphere, and Citrix XenServer.

### 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. The card enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1385 card supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure. Additional features and capabilities include:

- Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect
- The Cisco UCS VIC 1385 Virtual Interface Card provides high network performance and low latency for the most demanding applications such as SMB-Direct, VMQ, DPDK, and Cisco NetFlow

### 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. The mLOM card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1227T Virtual Interface Card supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment. Additional features and capabilities include:

- Stateless and agile design - The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure.
- Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect.
- Cisco SingleConnect technology provides an exceptionally easy, intelligent, and efficient way to connect and manage computing in the data center. Cisco SingleConnect technology dramatically simplifies the way that data centers connect to rack and blade servers, physical servers, virtual machines, LANs, SANs, and management networks.

### 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. The card enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1387 card supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure. Additional features and capabilities include:

- Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect
- The Cisco UCS VIC 1387 Virtual Interface Card provides high network performance and low latency for the most demanding applications such as SMB-Direct, VMQ, DPDK, and Cisco NetFlow

## Viewing Network Adapter Properties

### Before You Begin

- The server must be powered on, or the properties will not display.

### Procedure

- 
- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Network Adapters** tab.
- Step 4** In the **Network Adapters** area, review the following information:

Name	Description
Slot ID column	The slot in which the adapter is installed.
Product Name column	The product name for the adapter.
Number of Interfaces column	The number of interfaces for the adapter.
External Ethernet Interfaces	<b>ID</b> —The ID for the external ethernet interface. <b>MAC Address</b> —The MAC address for the external ethernet interface.

**Step 5** In the **Adapter Card** area, review the following information:

Name	Description
<b>Slot</b> column	The slot in which the network adapter resides.
<b>Product Name</b> column	The product name of the network adapter.
<b>Number of Interfaces</b> column	The number of interfaces for the network adapter.
<b>External Ethernet Interfaces</b> column	
<b>ID</b> column	The ID number of the external ethernet interface.
<b>MAC Address</b> column	The MAC address of the external ethernet interface.

## Viewing VIC Adapter Properties

### Before You Begin

- The server must be powered on, or the properties will not display.
- A supported Virtual Interface Card (VIC) must be installed in the chassis and the server must be powered on.

### Procedure

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

**Step 2** On the **Server** tab, click **Inventory**.

**Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.

**Step 4** In the **Adapter Cards** area, click an adapter in the table to display its properties.  
The resources of the selected adapter appear in the tabbed menu below the **Adapter Cards** area.

**Step 5** In the **Adapter Cards** area, review the following information for the installed adapters:

Name	Description
<b>PCI Slot</b> column	The PCI slot in which the adapter is installed.
<b>Product Name</b> column	The product name for the adapter.
<b>Serial Number</b> column	The serial number for the adapter.

Name	Description
<b>Product ID</b> column	The product ID for the adapter.
<b>Vendor</b> column	The vendor for the adapter.
<b>Cisco IMC Management Enabled</b> column	Whether the adapter is able to manage Cisco IMC. This functionality depends on the type of adapter installed and how it is configured. For details, see the hardware installation guide for the type of server you are using.

**Step 6** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.

**Step 7** In the **Adapter Card Properties** area, review the following information for the adapter:

Name	Description
<b>PCI Slot</b> field	The PCI slot in which the adapter is installed. <b>Note</b> For the C220 M4 and C240 M4 servers, PCI slot could also display as <b>MLOM</b> .
<b>Vendor</b> field	The vendor for the adapter.
<b>Product Name</b> field	The product name for the adapter.
<b>Product ID</b> field	The product ID for the adapter.
<b>Serial Number</b> field	The serial number for the adapter.
<b>Version ID</b> field	The version ID for the adapter.
<b>Hardware Revision</b> field	The hardware revision for the adapter.
<b>Cisco IMC Management Enabled</b> field	If this field displays <b>yes</b> , then the adapter is functioning in Cisco Card Mode and passing Cisco IMC management traffic through to the server Cisco IMC.
<b>Configuration Pending</b> field	If this field displays <b>yes</b> , 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.
<b>Description</b> field	The user-defined description for the adapter, if any.
<b>FIP Mode</b> field	Whether FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards.
<b>LLDP</b> field	Whether the LLDP option is enabled for this VIC card.  <b>Note</b> This option is available only on some UCS C-Series servers.

Name	Description
<b>VNTAG Mode</b> field	Whether virtual network tag (VNTAG) is enabled. If VNTAG mode is enabled: <ul style="list-style-type: none"> <li>• vNICs and vHBAs can be assigned to a specific channel</li> <li>• vNICs and vHBAs can be associated with a port profile</li> <li>• vNICs can fail over to another vNIC if there are communication problems</li> </ul>
<b>ISCSI Boot Capable</b> field	Whether iSCSI boot is supported on the adapter.
<b>usNIC Capable</b> field	Whether the adapter and the firmware running on the adapter support the usNIC.

**Step 8** In the **External Ethernet Interfaces** area, review the following information for the adapter:

Name	Description
<b>ID</b> column	The uplink port ID.
<b>MAC Address</b> column	The MAC address of the uplink port.
<b>Link State</b> column	The current operational state of the uplink port. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Fault</b></li> <li>• <b>Link Up</b></li> <li>• <b>Link Down</b></li> <li>• <b>SFP ID Error</b></li> <li>• <b>SFP Not Installed</b></li> <li>• <b>SFP Security Check Failed</b></li> <li>• <b>Unsupported SFP</b></li> </ul>
<b>Encap</b> column	The mode in which adapter operates. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>CE</b>—Classical Ethernet mode.</li> <li>• <b>NIV</b>—Network Interface Virtualization mode.</li> </ul>

Name	Description
<b>Admin Speed</b> column	<p>The data transfer rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b></li> <li>• <b>1 Gpbs</b></li> <li>• <b>10 Gpbs</b></li> <li>• <b>40 Gpbs</b></li> </ul> <p><b>Note</b> This option is only available for some adapter cards.</p>
<b>Operating Speed</b> column	<p>The operating rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b></li> <li>• <b>1 Gpbs</b></li> <li>• <b>10 Gpbs</b></li> <li>• <b>40 Gpbs</b></li> </ul> <p><b>Note</b> This option is only available for some adapter cards.</p>
<b>Training Link</b> column	<p>Indicates if link training is enabled on the port.</p>
<b>Connector Present</b> column	<p>Indicated whether or not the connector is present. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—Connector is present.</li> <li>• <b>No</b>—Connector not present.</li> </ul> <p><b>Note</b> This option is only available for some adapter cards.</p>
<b>Connector Supported</b> 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"> <li>• <b>Yes</b>—The connector is supported by Cisco.</li> <li>• <b>No</b>—The connector is not supported by Cisco.</li> </ul> <p>If the connector is not supported then the link will not be up.</p> <p><b>Note</b> This option is only available for some adapter cards.</p>
<b>Connector Type</b> column	<p>The type of the connector.</p> <p><b>Note</b> This option is only available for some adapter cards.</p>

Name	Description
<b>Connector Vendor</b> column	The vendor for the connector. <b>Note</b> This option is only available for some adapter cards.
<b>Connector Part Number</b> column	The part number of the connector. <b>Note</b> This option is only available for some adapter cards.
<b>Connector Part Revision</b> column	The part revision number of the connector. <b>Note</b> This option is only available for some adapter cards.

**Step 9** In the **Firmware** area, review the following information for the adapter:

Name	Description
<b>Running Version</b> field	The firmware version that is currently active.
<b>Backup Version</b> field	The alternate firmware version installed on the adapter, if any. The backup version is not currently running. To activate it, administrators can click <b>Activate Firmware</b> in the <b>Actions</b> area. <b>Note</b> 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.
<b>Startup Version</b> field	The firmware version that will become active the next time the adapter is rebooted.
<b>Bootloader Version</b> field	The bootloader version associated with the adapter card.
<b>Status</b> field	The status of the last firmware activation that was performed on this adapter. <b>Note</b> The status is reset each time the adapter is rebooted.

### What to Do Next

To view the properties of virtual NICs and virtual HBAs, see the following sections:

- [Viewing vNIC Properties, on page 25](#)
- [Viewing vHBA Properties, on page 11](#)

# Viewing Storage Adapter Properties

## Before You Begin

- The server must be powered on.

## Procedure

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

**Step 2** On the **Server** tab, click **Inventory**.

**Step 3** In the **Inventory** pane, click **Storage Adapters** tab and review the following information:

Name	Description
<b>Controller</b> field	The type of controller.
<b>PCI Slot</b> field	The PCI slot in which the adapter is installed.
<b>Product Name</b> field	The product name for the adapter.
<b>Serial Number</b> field	The serial number for the adapter.
<b>Firmware Package Build</b> field	The installed firmware package for the adapter.
<b>Product ID</b> field	The product ID for the adapter.
<b>Battery Status</b> field	The vendor for the adapter.
<b>Cache Memory Size</b> field	The size of the cache memory, in megabytes.
<b>Health</b> field	The health of the adapter. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Good</b></li> <li>• <b>Moderate Fault</b></li> <li>• <b>Severe Fault</b></li> <li>• <b>N/A</b></li> </ul>
<b>Details</b> link	Click the <b>Details</b> link to view the <b>Storage</b> tab.

# Managing vHBAs

## Guidelines for Managing vHBAs

When managing vHBAs, consider the following guidelines and restrictions:

- The Cisco UCS P81E Virtual Interface Card and Cisco UCS VIC 1225 Virtual Interface Card provide two vHBAs (fc0 and fc1). You can create up to 16 additional vHBAs on these adapter cards.



**Note** If Network Interface Virtualization (NIV) mode is enabled for the adapter, you must assign a channel number to a vHBA when you create it.

- When using the Cisco UCS P81E Virtual Interface Card or Cisco UCS VIC 1225 Virtual Interface Card 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

### Procedure

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Properties** to open the **vHBA Properties** dialog box.
- Step 8** In the **General** area, 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.
Target WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select <b>AUTO</b> . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.

Name	Description
<b>Target WWPN</b> field	The WWPN associated with the vHBA. To let the system generate the WWPN, select <b>AUTO</b> . To specify a WWPN, click the second radio button and enter the WWPN in the corresponding field.
<b>FC SAN Boot</b> check box	If checked, the vHBA can be used to perform a SAN boot.
<b>Enable Persistent LUN Binding</b> check box	If checked, any LUN ID associations are retained in memory until they are manually cleared.
<b>Uplink Port</b> field	The uplink port associated with the vHBA. <b>Note</b> This value cannot be changed for the system-defined vHBAs fc0 and fc1.
<b>MAC Address</b> field	The MAC address associated with the vHBA. To let the system generate the MAC address, select <b>AUTO</b> . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
<b>Default VLAN</b> field	If there is no default VLAN for this vHBA, click <b>NONE</b> . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
<b>Class of Service</b> drop-down list	The CoS for the vHBA. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. <b>Note</b> This option cannot be used in VNTAG mode.
<b>Rate Limit</b> field	The data rate limit for traffic on this vHBA, in Mbps. If you want this vHBA to have an unlimited data rate, select <b>OFF</b> . Otherwise, click the second radio button and enter an integer between 1 and 10,000. <b>Note</b> This option cannot be used in VNTAG mode.
<b>PCIe Device Order</b> field	The order in which this vHBA will be used. To let the system set the order, select <b>ANY</b> . To specify an order, select the second radio button and enter an integer between 0 and 17.
<b>EDTOV</b> field	The error detect timeout value (EDTOV), which is the number of milliseconds to wait before the system assumes that an error has occurred. Enter an integer between 1,000 and 100,000. The default is 2,000 milliseconds.

Name	Description
<b>RATOV field</b>	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.  Enter an integer between 5,000 and 100,000. The default is 10,000 milliseconds.
<b>Max Data Field Size field</b>	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports.  Enter an integer between 256 and 2112.
<b>Channel Number field</b>	The channel number that will be assigned to this vHBA.  Enter an integer between 1 and 1,000.  <b>Note</b> VNTAG mode is required for this option.
<b>Port Profile drop-down list</b>	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.  <b>Note</b> VNTAG mode is required for this option.

**Step 9** In the **Error Recovery** area, review the information in the following fields:

Name	Description
<b>Enable FCP Error Recovery</b> check box	If checked, the system uses FCP Sequence Level Error Recovery protocol (FC-TAPE).
<b>Link Down Timeout field</b>	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.
<b>Port Down I/O Retries field</b>	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.
<b>I/O Timeout Retry field</b>	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.

Name	Description
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 10** In the **Fibre Channel Interrupt** area, review the information in 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"> <li>• <b>MSIx</b>—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option.</li> <li>• <b>MSI</b>—MSI only.</li> <li>• <b>INTx</b>—PCI INTx interrupts.</li> </ul>

**Step 11** 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 1,024. The recommended value is 1024.
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 12** 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 <b>INFINITE</b> radio button. Otherwise select the second radio button and enter an integer into the corresponding field.

Name	Description
<b>FLOGI Timeout</b> 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 13** In the **Fibre Channel Port PLOGI** area, review the information in the following fields:

Name	Description
<b>PLOGI Retries</b> 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.
<b>PLOGI Timeout</b> 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 14** In the **SCSI I/O** area, review the information in the following fields:

Name	Description
<b>CDB Transmit Queue Count</b> field	The number of SCSI I/O queue resources the system should allocate. Enter an integer between 1 and 8.
<b>CDB Transmit Queue Ring Size</b> field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

**Step 15** In the **Receive/Transmit Queues** area, review the information in the following fields:

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

## Modifying vHBA Properties

### Procedure

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Properties** to open the **vHBA Properties** dialog box.
- Step 8** 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.
Target WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select <b>AUTO</b> . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
Target WWPN field	The WWPN associated with the vHBA. To let the system generate the WWPN, select <b>AUTO</b> . 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.
Enable Persistent LUN Binding check box	If checked, any LUN ID associations are retained in memory until they are manually cleared.
Uplink Port field	The uplink port associated with the vHBA. <b>Note</b> 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 <b>AUTO</b> . To specify an address, click the second radio button and enter the MAC address in the corresponding field.

Name	Description
<b>Default VLAN</b> field	If there is no default VLAN for this vHBA, click <b>NONE</b> . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
<b>Class of Service</b> drop-down list	<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><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>Rate Limit</b> 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 <b>OFF</b>. Otherwise, click the second radio button and enter an integer between 1 and 10,000.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>PCIe Device Order</b> field	<p>The order in which this vHBA will be used.</p> <p>To let the system set the order, select <b>ANY</b>. To specify an order, select the second radio button and enter an integer between 0 and 17.</p>
<b>EDTOV</b> 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>
<b>RATOV</b> 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>
<b>Max Data Field Size</b> 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>
<b>Channel Number</b> field	<p>The channel number that will be assigned to this vHBA.</p> <p>Enter an integer between 1 and 1,000.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>

Name	Description
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><b>Note</b> VNTAG mode is required for this option.</p>

**Step 9** In the **Error Recovery** area, update the following fields:

Name	Description
Enable 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 Retries 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>
I/O 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 10** In the **Fibre Channel Interrupt** area, update 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"> <li>• <b>MSIx</b>—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option.</li> <li>• <b>MSI</b>—MSI only.</li> <li>• <b>INTx</b>—PCI INTx interrupts.</li> </ul>

**Step 11** In the **Fibre Channel Port** area, update the following fields:

Name	Description
<b>I/O Throttle Count</b> 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.
<b>LUNs per Target</b> 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 1,024. The recommended value is 1024.
<b>LUN Queue Depth</b> 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 12** In the **Fibre Channel Port FLOGI** area, update the following fields:

Name	Description
<b>FLOGI Retries</b> 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 <b>INFINITE</b> radio button. Otherwise select the second radio button and enter an integer into the corresponding field.
<b>FLOGI Timeout</b> 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 13** In the **Fibre Channel Port PLOGI** area, update the following fields:

Name	Description
<b>PLOGI Retries</b> 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.
<b>PLOGI Timeout</b> 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 14** In the **SCSI I/O** area, update the following fields:

Name	Description
<b>CDB Transmit Queue Count</b> field	The number of SCSI I/O queue resources the system should allocate. Enter an integer between 1 and 8.
<b>CDB Transmit Queue Ring Size</b> field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

**Step 15** In the **Receive/Transmit Queues** area, update the following fields:

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

**Step 16** Click **Save Changes**.

---

## Creating a vHBA

The adapter provides two permanent vHBAs. If NIV mode is enabled, you can create up to 16 additional vHBAs.

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, choose one of these actions:
  - To create a vHBA using default configuration settings, click **Add**.
  - To create a vHBA using the same configuration settings as an existing vHBA, select that vHBA and click **Clone**.

The **Add vHBA** dialog box appears.

- Step 7** In the **Add vHBA** dialog box, enter a name for the vHBA in the **Name** entry box.
- Step 8** Click **Add vHBA**.
- 

### What to Do Next

- Reboot the server to create the vHBA.
- If configuration changes are required, configure the new vHBA as described in [Modifying vHBA Properties](#), on page 16.

## Deleting a vHBA

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.  
**Note** You cannot delete either of the two default vHBAs, **fc0** or **fc1**.
- Step 7** Click **Delete** and click **OK** to confirm.
- 

## 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

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.

If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.

- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Boot Table** to open the **Boot Table** dialog box for the selected vHBA.
- Step 8** In the **Boot Table** dialog box, click **Add** to open the **Add Boot Entry** dialog box.
- Step 9** In the **Add Boot Entry** dialog box, update the following fields:

Name	Description
<b>Target WWPN</b> 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.
<b>LUN ID</b> field	The LUN ID that corresponds to the location of the boot image. Enter an ID between 0 and 255.
<b>Add Boot Entry</b> button	Adds the specified location to the boot table.
<b>Reset Values</b> button	Clears the values currently entered in the fields.
<b>Cancel</b> button	Closes the dialog box without saving any changes made while the dialog box was open.

- Step 10** Click **Add Boot Entry**.
- 

## Deleting a Boot Table Entry

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.

- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Boot Table** to open the **Boot Table** dialog box for the selected vHBA.
- Step 8** In the **Boot Table** dialog box, click the entry to be deleted.
- Step 9** Click **Delete** 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

### Procedure

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Persistent Bindings** to open the **Persistent Bindings** dialog box for the selected vHBA.
- Step 8** In the **Persistent Bindings** dialog box for the selected vHBA, review the following information:

Name	Description
<b>Index</b> column	The unique identifier for the binding.
<b>Target WWPN</b> column	The target World Wide Port Name with which the binding is associated.
<b>Host WWPN</b> column	The host World Wide Port Name with which the binding is associated.
<b>Bus ID</b> column	The bus ID with which the binding is associated.
<b>Target ID</b> column	The target ID on the host system with which the binding is associated.
<b>Rebuild Persistent Bindings</b> button	Clears all unused bindings and resets the ones that are in use.
<b>Close</b> button	Closes the dialog box and saves your changes.

**Step 9** Click **Close**.

---

## Rebuilding Persistent Bindings

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vHBAs** tab.
- Step 6** In the **Host Fibre Channel Interfaces** area, select a vHBA from the table.
- Step 7** Click **Persistent Bindings** to open the **Persistent Bindings** dialog box for the selected vHBA.
- Step 8** In the **Persistent Bindings** dialog box for the selected vHBA, click **Rebuild Persistent Bindings**.
- Step 9** Click **Close**.
- 

## Managing vNICs

### Guidelines for Managing vNICs

When managing vNICs, consider the following guidelines and restrictions:

- The Cisco UCS P81E Virtual Interface Card and Cisco UCS VIC 1225 Virtual Interface Card provide two default vNICs (eth0 and eth1). You can create up to 16 additional vNICs on these adapter cards.



**Note** If Network Interface Virtualization (NIV) 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.

Cisco C-series servers use Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) for packet transfers. RoCE defines the mechanism of performing RDMA over ethernet, based on the similar mechanism of RDMA over Infiniband. However, RoCE, with its performance oriented characteristics, delivers a superior performance compared to traditional network socket implementation because of the lower latency,

lower CPU utilization and higher utilization of network bandwidth. RoCE meets the requirement of moving large amount of data across networks very efficiently.

The RoCE firmware requires the following configuration parameters provided by Cisco UCS Manager for better vNIC performance:

- Queue Pairs
- Memory Regions
- Resource Groups

## Viewing vNIC Properties

### Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Ethernet Interfaces** pane's **vNIC Properties** area, review the information in 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. <b>Note</b> This feature works only when the <b>CDN Support for VIC</b> 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 <b>Auto</b> . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Class of Service drop-down list	The class of service to associate with traffic from this vNIC. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. <b>Note</b> This option cannot be used in VNTAG mode.

Name	Description
<b>Trust Host CoS</b> check box	Check this box if you want the vNIC to use the class of service provided by the host operating system.
<b>PCI Link</b> field	<p>The link through which vNICs can be connected. These are the following values:</p> <ul style="list-style-type: none"> <li>• <b>0</b> - The first cross-edged link where the vNIC is placed.</li> <li>• <b>1</b> - The second cross-edged link where the vNIC is placed.</li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• This option is available only on some Cisco UCS C-Series servers.</li> <li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li> </ul>
<b>PCI Order</b> field	<p>The order in which this vNIC will be used.</p> <p>To let the system set the order, select <b>Any</b>. To specify an order, select the second radio button and enter an integer between 0 and 17.</p>
<b>Default VLAN</b> field	<p>If there is no default VLAN for this vNIC, click <b>NONE</b>. Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>VLAN Mode</b> drop-down list	<p>If you want to use VLAN trunking, select <b>TRUNK</b>. Otherwise, select <b>ACCESS</b>. When the VLAN is set to <b>ACCESS</b> mode, any frame received from the specified default VLAN (1-4094) that is received from the switch with a TAG removes that TAG when it is sent to the host OS through the vNIC.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>Rate Limit</b> field	<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 or 40,000 Mbps depending on the adapter card you choose.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>Enable PXE Boot</b> check box	Check this box if the vNIC can be used to perform a PXE boot.
<b>Channel Number</b> field	<p>Select the channel number that will be assigned to this vNIC.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>

Name	Description
<b>Port Profile</b> drop-down list	<p>Select the port profile that should be associated with the vNIC.</p> <p>This field displays the port profiles defined on the switch to which this server is connected.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>
<b>Enable Uplink Failover</b> check box	<p>Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>
<b>Enable VMQ</b> check box	<p>Check this box to enable Virtual Machine Queue (VMQ).</p> <p><b>Note</b> Ensure that VMQ is not enabled when SR-IOV or netflow option is enabled on the adapter.</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>
<b>Enable aRFS</b> check box	<p>Check this box to enable Accelerated Receive Flow steering (aRFS).</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>
<b>Enable NVGRE</b> check box	<p>Check this box to enable Network Virtualization using Generic Routing Encapsulation.</p> <ul style="list-style-type: none"> <li>• This option is available only on some Cisco UCS C-Series servers.</li> <li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li> </ul>
<b>Enable VXLAN</b> check box	<p>Check this box to enable Virtual Extensible LAN.</p> <ul style="list-style-type: none"> <li>• This option is available only on some Cisco UCS C-Series servers.</li> <li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li> </ul>
<b>Failback Timeout</b> field	<p>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.</p> <p>Enter a number of seconds between 0 and 600.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>

**Step 5** In the **Ethernet Interrupt** area, review the information in the following fields:

Name	Description
<b>Interrupt Count</b> 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.
<b>Coalescing Time</b> 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.
<b>Coalescing Type</b> drop-down list	This can be one of the following: <ul style="list-style-type: none"> <li>• <b>MIN</b>—The system waits for the time specified in the <b>Coalescing Time</b> field before sending another interrupt event.</li> <li>• <b>IDLE</b>—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the <b>Coalescing Time</b> field.</li> </ul>
<b>Interrupt Mode</b> drop-down list	The preferred driver interrupt mode. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>MSI-X</b>—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option.</li> <li>• <b>MSI</b>—MSI only.</li> <li>• <b>INTx</b>—PCI INTx interrupts.</li> </ul>

**Step 6** In the **Ethernet Receive Queue** area, review the information in the following fields:

Name	Description
<b>Receive Queue Count</b> field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
<b>Receive Queue Ring Size</b> field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.

**Step 7** In the **Ethernet Transmit Queue** area, review the information in the following fields:

Name	Description
<b>Transmit Queue Count</b> field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.

Name	Description
Transmit Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.

**Step 8** In the **Completion Queue** area, review the information in the following fields:

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

**Step 9** In the **TCP Offload** area, review the information in the following fields:

Name	Description
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. <b>Note</b> 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.
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.

**Step 10** In the **Receive Side Scaling** area, review the information in the following fields:

Name	Description
<b>Enable TCP Receive Side Scaling</b> 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.
<b>Enable IPv4 RSS</b> check box	If checked, RSS is enabled on IPv4 networks.
<b>Enable TCP-IPv4 RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv4 networks.
<b>Enable IPv6 RSS</b> check box	If checked, RSS is enabled on IPv6 networks.
<b>Enable TCP-IPv6 RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.
<b>Enable IPv6 Extension RSS</b> check box	If checked, RSS is enabled for IPv6 extensions.
<b>Enable TCP-IPv6 Extension RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.

## Modifying vNIC Properties

### Procedure

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 6** In the **Host Ethernet Interfaces** area, select a vNIC from the table.
- Step 7** Click **Properties** to open the **vNIC Properties** dialog box.
- Step 8** In the **General** area, 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. <b>Note</b> This feature works only when the <b>CDN Support for VIC</b> 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 <b>Auto</b> . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Class of Service drop-down list	The class of service to associate with traffic from this vNIC. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. <b>Note</b> This option cannot be used in VNTAG mode.
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 Link field	The link through which vNICs can be connected. These are the following values: <ul style="list-style-type: none"><li>• <b>0</b> - The first cross-edged link where the vNIC is placed.</li><li>• <b>1</b> - The second cross-edged link where the vNIC is placed.</li></ul> <b>Note</b> <ul style="list-style-type: none"><li>• This option is available only on some Cisco UCS C-Series servers.</li><li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li></ul>
PCI Order field	The order in which this vNIC will be used. To let the system set the order, select <b>Any</b> . To specify an order, select the second radio button and enter an integer between 0 and 17.

Name	Description
<b>Default VLAN</b> field	<p>If there is no default VLAN for this vNIC, click <b>NONE</b>. Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>VLAN Mode</b> drop-down list	<p>If you want to use VLAN trunking, select <b>TRUNK</b>. Otherwise, select <b>ACCESS</b>. When the VLAN is set to <b>ACCESS</b> mode, any frame received from the specified default VLAN (1-4094) that is received from the switch with a TAG removes that TAG when it is sent to the host OS through the vNIC.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>Rate Limit</b> field	<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 or 40,000 Mbps depending on the adapter card you choose.</p> <p><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>Enable PXE Boot</b> check box	<p>Check this box if the vNIC can be used to perform a PXE boot.</p>
<b>Channel Number</b> field	<p>Select the channel number that will be assigned to this vNIC.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>
<b>Port Profile</b> drop-down list	<p>Select the port profile that should be associated with the vNIC.</p> <p>This field displays the port profiles defined on the switch to which this server is connected.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>
<b>Enable Uplink Failover</b> check box	<p>Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>
<b>Enable VMQ</b> check box	<p>Check this box to enable Virtual Machine Queue (VMQ).</p> <p><b>Note</b> Ensure that VMQ is not enabled when SR-IOV or netflow option is enabled on the adapter.</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>
<b>Enable aRFS</b> check box	<p>Check this box to enable Accelerated Receive Flow steering (aRFS).</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>

Name	Description
<p><b>Enable NVGRE</b> check box</p>	<p>Check this box to enable Network Virtualization using Generic Routing Encapsulation.</p> <ul style="list-style-type: none"> <li>• This option is available only on some Cisco UCS C-Series servers.</li> <li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li> </ul>
<p><b>Enable VXLAN</b> check box</p>	<p>Check this box to enable Virtual Extensible LAN.</p> <ul style="list-style-type: none"> <li>• This option is available only on some Cisco UCS C-Series servers.</li> <li>• This option is available only on C-Series servers with Cisco VIC 1385 cards.</li> </ul>
<p><b>Failback Timeout</b> field</p>	<p>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.</p> <p>Enter a number of seconds between 0 and 600.</p> <p><b>Note</b> VNTAG mode is required for this option.</p>

**Step 9** In the **Ethernet Interrupt** area, update the following fields:

Name	Description
<p><b>Interrupt Count</b> field</p>	<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>
<p><b>Coalescing Time</b> field</p>	<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>
<p><b>Coalescing Type</b> drop-down list</p>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>MIN</b>—The system waits for the time specified in the <b>Coalescing Time</b> field before sending another interrupt event.</li> <li>• <b>IDLE</b>—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 <b>Coalescing Time</b> field.</li> </ul>

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

**Step 10** In the **Ethernet Receive Queue** area, update the following fields:

Name	Description
<b>Receive Queue Count</b> field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
<b>Receive Queue Ring Size</b> field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.

**Step 11** In the **Ethernet Transmit Queue** area, update the following fields:

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

**Step 12** In the **Completion Queue** area, update the following fields:

Name	Description
<b>Completion Queue Count</b> 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.
<b>Completion Queue Ring Size</b> field	The number of descriptors in each completion queue. This value cannot be changed.

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

Name	Description
RoCE checkbox	Check the check box to change the RoCE Properties.
Queue Pairs (1 - 8192) field	<p>The number of queue pairs per adapter. Enter an integer between 1 and 8192.</p> <p>We recommend that this number be an integer power of 2. The recommended value for queue pairs per vNIC is 2048. This allows four vNICs to be created per adapter. Windows driver reserves two queue pairs for internal use, so a valid range of values would be 4 to 8192 queue pairs per vNIC.</p>
Memory Regions (1 - 524288) field	<p>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.</p> <p>The number of memory regions supported should be enough to meet application requirements as the regions are primarily used to send operation channel semantics.</p>
Resource Groups (1 - 128) field	<p>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. Recommended value is 32.</p> <p>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.</p>

**Step 14** In the **TCP Offload** area, update the following fields:

Name	Description
<p><b>Enable TCP Segmentation Offload</b> check box</p>	<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><b>Note</b> This option is also known as Large Send Offload (LSO).</p>
<p><b>Enable TCP Rx Offload Checksum Validation</b> check box</p>	<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>

Name	Description
<b>Enable TCP Tx Offload Checksum Generation</b> 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.
<b>Enable Large Receive</b> 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.

**Step 15** In the **Receive Side Scaling** area, update the following fields:

Name	Description
<b>Enable TCP Receive Side Scaling</b> 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.
<b>Enable IPv4 RSS</b> check box	If checked, RSS is enabled on IPv4 networks.
<b>Enable TCP-IPv4 RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv4 networks.
<b>Enable IPv6 RSS</b> check box	If checked, RSS is enabled on IPv6 networks.
<b>Enable TCP-IPv6 RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.
<b>Enable IPv6 Extension RSS</b> check box	If checked, RSS is enabled for IPv6 extensions.
<b>Enable TCP-IPv6 Extension RSS</b> check box	If checked, RSS is enabled for TCP transmissions across IPv6 networks.

**Step 16** Click **Save Changes**.

---

## Creating a vNIC

The adapter provides two permanent vNICs. You can create up to 16 additional vNICs.

## Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 6** In the **Host Ethernet Interfaces** area, choose one of these actions:
- To create a vNIC using default configuration settings, click **Add**.
  - To create a vNIC using the same configuration settings as an existing vNIC, select that vNIC and click **Clone**.

The **Add vNIC** dialog box appears.

- Step 7** In the **Add vNIC** dialog box, enter a name for the vNIC in the **Name** entry box.
- Step 8** (Optional) In the **Add vNIC** dialog box, enter a channel number for the vNIC in the **Channel Number** entry box.
- Note** If NIV is enabled on the adapter, you must assign a channel number for the vNIC when you create it.
- Step 9** Click **Add vNIC**.
- 

## What to Do Next

If configuration changes are required, configure the new vNIC as described in [Modifying vNIC Properties](#), on page 30.

# Deleting a vNIC

## Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 6** 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 7** Click **Delete** and click **OK** to confirm.

---

## 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.

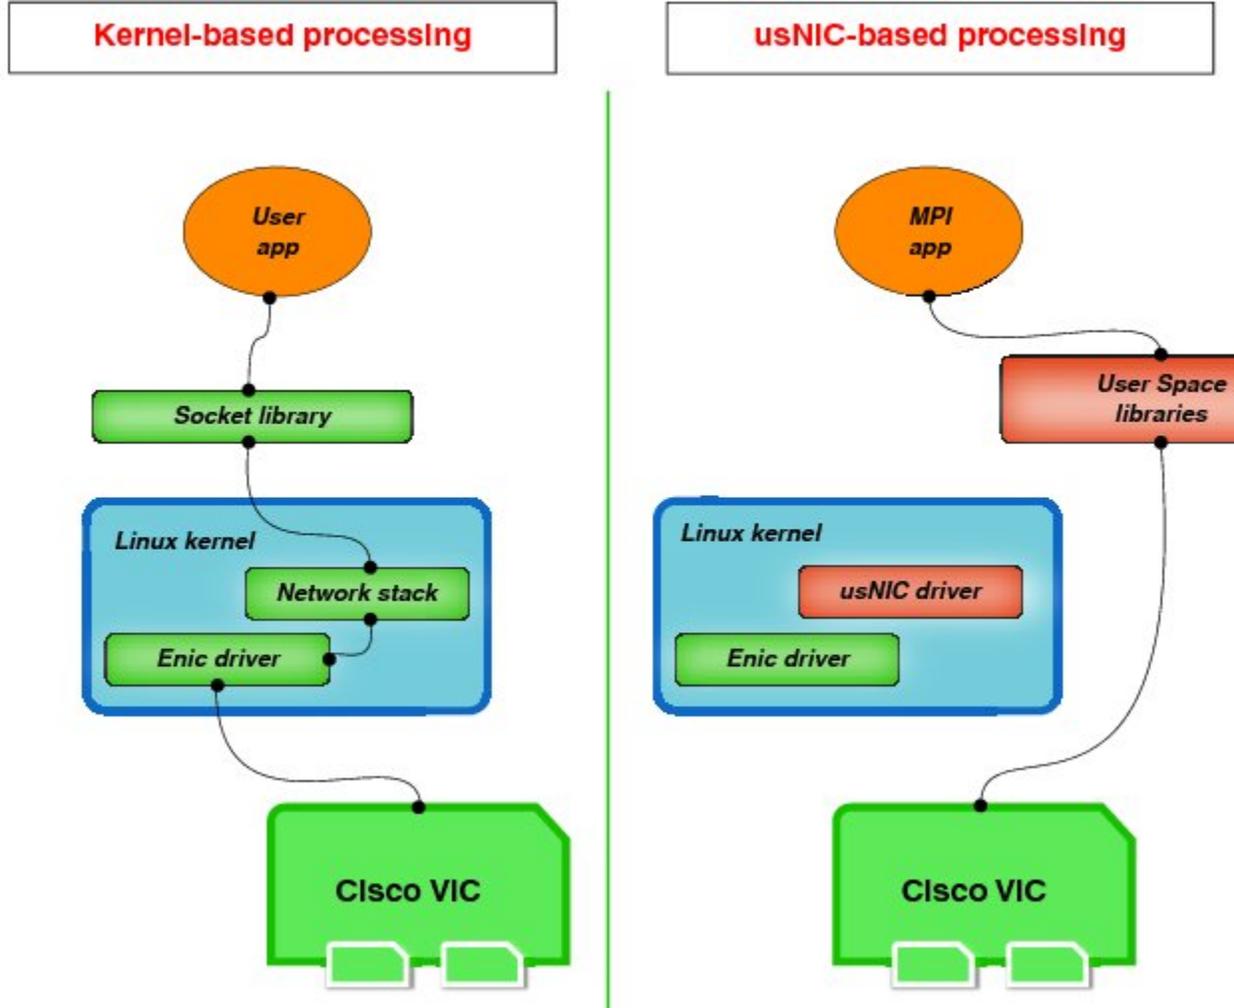
Cisco usNIC offers the following benefits for your MPI applications:

- 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



## Configuring Cisco usNIC Using the Cisco IMC GUI



**Note** Even though several properties are listed for Cisco usNIC in the usNIC properties dialog box, you must configure only the following properties because the other properties are not currently being used.

- **cq-count**
- **rq-count**
- **tq-count**
- **usnic-count**

### Before You Begin

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

### Procedure

- 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 **Server** tab.
- Step 3** On the **Server** tab, click **Inventory**.
- Step 4** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 5** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 6** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 7** In the **Host Ethernet Interfaces** area, select a vNIC from the table.
- Note** For each vNIC that you want to configure as a usNIC, select the vNIC entry from the table and specify its properties as explained in steps 9 through step 18.
- Step 8** Click **usNIC** to open the **usNIC Properties** dialog box.
- Step 9** In the **usNICs** property, specify the number of Cisco usNICs that you want to create.  
Each MPI process that is running on the server requires a dedicated usNIC. You might need to create up to 64 usNICs to sustain 64 MPI processes running simultaneously. We recommend that you create at least as many usNICs, per usNIC-enabled vNIC, as the number of physical cores on your server. For example, if you have 8 physical cores on your server, create 8 usNICs.
- Step 10** In the **Properties** area, update the following fields:

Field Name	Description
<b>Transmit Queue Count</b>	The number of transmit queue resources to allocate. Cisco recommends setting this value to 6.

Field Name	Description
Receive Queue Count	The number of receive queue resources to allocate. Cisco recommends setting this value to 6.
Completion Queue Count	The number of completion queue resources to allocate. Cisco recommends setting this value to 6.

**Step 11** Click **Apply**.

**Step 12** In the **Navigation** pane, click the **Server** tab.

**Step 13** On the **Server** tab, click **BIOS**.

**Step 14** In the **Actions** area, click **Configure BIOS**.

**Step 15** In the **Configure BIOS Parameters** dialog box, click the **Advanced** tab.

**Step 16** In the **Processor Configuration** area, set the following properties to Enabled:

- **Intel(R) VT-d**
- **Intel(R) VT-d ATS support**
- **Intel(R) VT-d Coherency Support**

**Step 17** Click **Save Changes**.

The changes take effect upon the next server reboot.

## Viewing usNIC Properties

### Procedure

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

**Step 2** On the **Server** tab, click **Inventory**.

**Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.

**Step 4** In the **Adapter Cards** area, select the adapter card.

If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.

**Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.

**Step 6** In the **Host Ethernet Interface** area, select the usNIC that is assigned to vNIC, to open the **usNIC properties** dialog box.

**Step 7** In the **usNIC** area, review or update the information in the following fields:

Name	Description
Name	The name for the vNIC that is the parent of the usNIC. <b>Note</b> 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.

**Step 8** In the **Properties** area, review or update the information in the following fields:

Name	Description
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.

Name	Description
<b>Interrupt Coalescing Type</b> drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>MIN</b>—The system waits for the time specified in the <b>Coalescing Time</b> field before sending another interrupt event.</li> <li>• <b>IDLE</b>—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the <b>Coalescing Time</b> field.</li> </ul>
<b>Interrupt Coalescing Timer Time</b> 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>
<b>Class of Service</b> 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><b>Note</b> This option cannot be used in VNTAG mode.</p>
<b>TCP Segment Offload</b> 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><b>Note</b> This option is also known as Large Send Offload (LSO).</p>
<b>Large Receive</b> 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>
<b>TCP Tx Checksum</b> 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>
<b>TCP Rx Checksum</b> 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>

Name	Description
Apply button	Applies changes to all the usNICs associated with the vNIC device.
Reset values button	Restores the values for the usNIC to the settings that were in effect when this dialog box was first opened.
Cancel button	Closes the dialog box without making any changes.

## Configuring iSCSI Boot Capability

### Configuring iSCSI Boot Capability for vNICs

When the rack-servers are configured in a standalone mode, and when the VIC adapters are directly attached to the Nexus 5000 family of switches, you can configure these VIC adapters to boot the servers remotely from iSCSI storage targets. You can configure Ethernet vNICs to enable a rack server to load the host OS image from remote iSCSI target devices.

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

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

#### Before You Begin

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

**Procedure**

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 6** In the **Host Ethernet Interfaces** area, select a vNIC from the table, and click **iSCSI Boot**.
- Step 7** 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.  <b>Note</b> 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)
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 8** In the **Initiator Area**, update the following fields:

Name	Description
Name field	<p>A regular expression that defines the name of the iSCSI initiator.</p> <p>You can enter any alphanumeric string as well as the following special characters:</p> <ul style="list-style-type: none"> <li>• . (period)</li> <li>• : (colon)</li> <li>• - (dash)</li> </ul> <p><b>Note</b> 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.
Secondary DNS field	The secondary DNS server address.
TCP Timeout field	<p>The number of seconds to wait before the initiator assumes that TCP is unavailable.</p> <p>Enter an integer between 0 and 255 (default: 15 seconds)</p>
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 **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.

Name	Description
CHAP Secret field	The Challenge-Handshake Authentication Protocol (CHAP) shared secret of the initiator.

**Step 10** 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 11** Click **Configure iSCSI**.

## Removing iSCSI Boot Configuration from a vNIC

### Before You Begin

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

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 6** In the **Host Ethernet Interfaces** area, select a vNIC from the table, and click **iSCSI Boot**.
- Step 7** In the dialog box that appears, click **Unconfigure iSCSI**.
- 

## Configuring Virtual Machine Queues on a vNIC

### Before You Begin

You must log in to the Cisco IMC GUI with administrator privileges to perform this task.

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the **Ethernet Interfaces** pane's **vNIC Properties** area, check the **Enable VMQ** checkbox.
- Step 4** In the **Ethernet Transmit Queue** area, enter an integer in the **Transmit Queue Count** field. This number should be greater than 1.
- Step 5** In the **Ethernet Receive Queue** area, enter an integer in the **Receive Queue Count** field. This number should be equal to the number of transmit queues.
- Step 6** In the **Ethernet Interrupt** area, enter an integer in the **Interrupt Count** field. This should be equal to the number of logical processors, or completion queues.
- 

### What to Do Next

- Reboot the server.
- Create a logical switch on the NIC.

# 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.

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
- Step 6** In the **Actions** area of the **General** tab, click **Export Configuration**.  
The **Export Adapter Configuration** dialog box opens.
- Step 7** In the **Export Adapter Configuration** dialog box, update the following fields:

Name	Description
<b>Export to</b> drop-down list	<p>The remote server type. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>TFTP Server</b></li> <li>• <b>FTP Server</b></li> <li>• <b>SFTP Server</b></li> <li>• <b>SCP Server</b></li> <li>• <b>HTTP Server</b></li> </ul> <p><b>Note</b> 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 &lt;server_finger_print_ID&gt; 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>
<b>Server IP/Hostname</b> field	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 <b>Export to</b> drop-down list, the name of the field may vary.
<b>Path and Filename</b> field	The path and filename Cisco IMC should use when exporting the file to the remote server.
<b>Username</b>	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.
<b>Password</b>	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.

**Step 8** Click **Export Configuration**.

---

## Importing the Adapter Configuration

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.

If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.

**Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.

**Step 6** In the **Actions** area of the **General** tab, click **Import Configuration**.  
The **Import Adapter Configuration** dialog box opens.

**Step 7** In the **Import Adapter Configuration** dialog box, update the following fields:

Name	Description
<b>Import from</b> drop-down list	The remote server type. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>TFTP Server</b></li> <li>• <b>FTP Server</b></li> <li>• <b>SFTP Server</b></li> <li>• <b>SCP Server</b></li> <li>• <b>HTTP Server</b></li> </ul> <p><b>Note</b> 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 &lt;server_finger_print_ID&gt; 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>
<b>Server IP/Hostname</b> field	The IPv4 or IPv6 address, or hostname of the server on which the adapter configuration file resides. Depending on the setting in the <b>Import from</b> drop-down list, the name of the field may vary.
<b>Path and Filename</b> field	The path and filename of the configuration file on the remote server.
<b>Username</b>	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.
<b>Password</b>	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.

**Step 8** Click **Import Configuration**.  
The adapter downloads the configuration file from the specified path on the TFTP server at the specified IP address. The configuration will be installed during the next server reboot.

**What to Do Next**

Reboot the server to apply the imported configuration.

## Restoring Adapter Defaults

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
- Step 6** In the **Actions** area of the **General** tab, click **Reset To Defaults** and click **OK** to confirm.
- 

## Managing Adapter Firmware

### Adapter Firmware

A Cisco UCS C-Series network adapter contains the following firmware components:

- Adapter firmware—The main operating firmware, consisting of an active and a backup image, can be installed from the Cisco IMC GUI or CLI interface or from the Host Upgrade Utility (HUU). You can upload a firmware image from either a local file system or a TFTP server.
- Bootloader firmware—The bootloader firmware cannot be installed from the Cisco IMC. You can install this firmware using the Host Upgrade Utility.

## Installing Adapter Firmware From a Local File

### Before You Begin

Store the adapter firmware file in the file system of the managing computer.

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.

- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
  - Step 6** In the **Actions** area of the **General** tab, click **Install Firmware** to open the **Install Adapter Firmware** dialog box.
  - Step 7** In the **Install Adapter Firmware** dialog box, select **Install from local file**, then click **Next**.
  - Step 8** Click **Browse...** and locate the adapter firmware file.
  - Step 9** Click **Install Firmware**.
- 

### What to Do Next

To activate the new firmware, see *Activating Adapter Firmware*.

## Installing Adapter Firmware From a Remote Server

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
- Step 2** On the **Server** tab, click **Inventory**.
- Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
- Step 6** In the **Actions** area of the **General** tab, click **Install Firmware** to open the **Install Adapter Firmware** dialog box.
- Step 7** In the **Install Adapter Firmware** dialog box, select **Install from Remote Server**, then click **Next**.
- Step 8** In the **Install Adapter Firmware** dialog box, update the following fields:

Name	Description
<b>Install from</b> drop-down list	<p>The remote server type. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>TFTP Server</b></li> <li>• <b>FTP Server</b></li> <li>• <b>SFTP Server</b></li> <li>• <b>SCP Server</b></li> <li>• <b>HTTP Server</b></li> </ul> <p><b>Note</b> 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 &lt;server_finger_print_ID&gt; 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>
<b>Server IP/Hostname</b> field	The IP address or hostname of the server on which the adapter configuration file resides. Depending on the setting in the <b>Install from</b> drop-down list, the name of the field may vary.
<b>Path and Filename</b> field	The path and filename of the configuration file on the remote server.
<b>Username</b>	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.
<b>Password</b>	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.
<b>Back</b> button	Click this button if you want to specify a local path for the firmware package.
<b>Install Firmware</b> button	Click this button to install the selected firmware package in the adapter's backup memory slot.
<b>Close</b> button	Click this button to close the wizard without making any changes to the firmware versions stored on the server.

### Step 9 Click **Install Firmware**.

#### What to Do Next

To activate the new firmware, see *Activating Adapter Firmware*.

## Activating Adapter Firmware

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
  - Step 2** On the **Server** tab, click **Inventory**.
  - Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
  - Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
  - Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
  - Step 6** In the **Actions** area of the **General** tab, click **Activate Firmware** to open the **Activate Adapter Firmware** dialog box.
  - Step 7** In the **Activate Adapter Firmware** dialog box, select the image to run the next time the firmware starts up.
  - Step 8** Click **Activate Adapter Firmware**.
- 

## Resetting the Adapter

### Procedure

---

- Step 1** In the **Navigation** pane, click the **Server** tab.
  - Step 2** On the **Server** tab, click **Inventory**.
  - Step 3** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
  - Step 4** In the **Adapter Cards** area, select the adapter card.  
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
  - Step 5** In the tabbed menu below the **Adapter Cards** area, click the **General** tab.
  - Step 6** In the **Actions** area of the **General** tab, click **Reset** and click **Yes** to confirm.  
**Note** Resetting the adapter also resets the host.
-

