



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 2](#)
- [Viewing VIC Adapter Properties, page 3](#)
- [Viewing Storage Adapter Properties, page 7](#)
- [Managing vHBAs, page 8](#)
- [Managing vNICs, page 21](#)
- [Managing VM FEX, page 42](#)
- [Managing Storage Adapters, page 46](#)
- [Backing Up and Restoring the Adapter Configuration, page 62](#)
- [Managing Adapter Firmware, page 65](#)
- [Resetting the Adapter, page 67](#)

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 VIC1225 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 VIC1225 Virtual Interface Card

The Cisco UCS VIC1225 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.

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.

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

Name	Description
ID column	The ID for the external ethernet interface.
MAC Address column	The MAC address for 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
PCI Slot column	The PCI slot in which the adapter is installed.
Product Name column	The product name for the adapter.
Serial Number column	The serial number for the adapter.
Product ID column	The product ID for the adapter.
Vendor column	The vendor for the adapter.
Cisco IMC Management Enabled 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
PCI Slot field	The PCI slot in which the adapter is installed.
Vendor field	The vendor for the adapter.
Product Name field	The product name for the adapter.
Product ID field	The product ID for the adapter.
Serial Number field	The serial number for the adapter.
Version ID field	The version ID for the adapter.
Hardware Revision field	The hardware revision for the adapter.
Cisco IMC Management Enabled field	If this field displays yes , then the adapter is functioning in Cisco Card Mode and passing Cisco IMC management traffic through to the server Cisco IMC.
Configuration Pending field	If this field displays yes , the adapter configuration has changed in Cisco IMC but these changes have not been communicated to the host operating system. To activate the changes, an administrator must reboot the adapter.
Description field	The user-defined description for the adapter, if any.
FIP Mode field	Whether FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards.

Name	Description
VNTAG Mode field	Whether virtual network tag (VNTAG) is enabled. If VNTAG mode is enabled: <ul style="list-style-type: none"> • vNICs and vHBAs can be assigned to a specific channel • vNICs and vHBAs can be associated with a port profile • vNICs can fail over to another vNIC if there are communication problems
ISCSI Boot Capable field	Whether iSCSI boot is supported on the adapter.
usNIC Capable 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
ID column	The uplink port ID.
MAC Address column	The MAC address of the uplink port.
Link State column	The current operational state of the uplink port. This can be one of the following: <ul style="list-style-type: none"> • Fault • Link Up • Link Down • SFP ID Error • SFP Not Installed • SFP Security Check Failed • Unsupported SFP
Encap column	The mode in which adapter operates. This can be one of the following: <ul style="list-style-type: none"> • CE—Classical Ethernet mode. • NIV—Network Interface Virtualization mode.

Name	Description
Admin Speed column	<p>The data transfer rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto • 1 Gpbs • 10 Gpbs <p>Note This option is only available for some adapter cards.</p>
Operating Speed column	<p>The operating rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto • 1 Gpbs • 10 Gpbs <p>Note This option is only available for some adapter cards.</p>

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

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

What to Do Next

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

- [Viewing vNIC Properties, on page 22](#)
- [Viewing Virtual FEX Properties, on page 43](#)
- [Viewing vHBA Properties, on page 8](#)

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
Controller field	The type of controller.
PCI Slot field	The PCI slot in which the adapter is installed.
Product Name field	The product name for the adapter.
Serial Number field	The serial number for the adapter.
Firmware Package Build field	The installed firmware package for the adapter.
Product ID field	The product ID for the adapter.
Battery Status field	The vendor for the adapter.
Cache Memory Size field	The size of the cache memory, in megabytes.
Health field	The health of the adapter. This can be one of the following: <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault • N/A
Details link	Click the Details link to view the Storage 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 VIC1225 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 VIC1225 Virtual Interface Card in an FCoE application, you must associate the vHBA with the FCoE VLAN. Follow the instructions in [Modifying vHBA Properties](#), on page 12 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.

Name	Description
World Wide Node Name field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
World Wide Port Name field	The WWPN associated with the vHBA. To let the system generate the WWPN, select AUTO . To specify a WWPN, click the second radio button and enter the WWPN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
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. Note This value cannot be changed for the system-defined vHBAs fc0 and fc1.
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
Class of Service 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. Note This option cannot be used in VNTAG mode.
Rate Limit field	The data rate limit for traffic on this vHBA, in Mbps. If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter an integer between 1 and 10,000. Note This option cannot be used in VNTAG mode.
PCIe Device Order field	The order in which this vHBA will be used. To let the system set the order, select ANY . To specify an order, select the second radio button and enter an integer between 0 and 17.

Name	Description
EDTOV 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.
RATOV field	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.
Max Data Field Size field	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports. Enter an integer between 256 and 2112.
Channel Number field	The channel number that will be assigned to this vHBA. Enter an integer between 1 and 1,000. Note VNTAG mode is required for this option.
Port Profile drop-down list	The port profile that should be associated with the vHBA, if any. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.

Step 9 In the **Error Recovery** area, review the information in 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	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240,000.
Port Down I/O Retries field	The number of times an I/O request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255.

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"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—MSI only. • INTx—PCI INTx interrupts.

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.

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 INFINITE radio button. Otherwise select the second radio button and enter an integer into the corresponding field.
FLOGI Timeout field	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1,000 and 255,000.

Step 13 In the **Fibre Channel Port PLOGI** area, review the information in the following fields:

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

Step 14 In the **SCSI I/O** area, review the information in the following fields:

Name	Description
CDB Transmit Queue Count field	The number of SCSI I/O queue resources the system should allocate. Enter an integer between 1 and 8.
CDB Work Queue Ring Size 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
FC Work Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 128.
FC Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 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.
World Wide Node Name field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
World Wide Port Name field	The WWPN associated with the vHBA. To let the system generate the WWPN, select AUTO . To specify a WWPN, click the second radio button and enter the WWPN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
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. Note This value cannot be changed for the system-defined vHBAs fc0 and fc1.
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
Class of Service 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. Note This option cannot be used in VNTAG mode.

Name	Description
Rate Limit field	The data rate limit for traffic on this vHBA, in Mbps. If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter an integer between 1 and 10,000. Note This option cannot be used in VNTAG mode.
PCIe Device Order field	The order in which this vHBA will be used. To let the system set the order, select ANY . To specify an order, select the second radio button and enter an integer between 0 and 17.
EDTOV 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.
RATOV field	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.
Max Data Field Size field	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports. Enter an integer between 256 and 2112.
Channel Number field	The channel number that will be assigned to this vHBA. Enter an integer between 1 and 1,000. Note VNTAG mode is required for this option.
Port Profile drop-down list	The port profile that should be associated with the vHBA, if any. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.

Step 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).

Name	Description
Link Down Timeout field	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240,000.
Port Down I/O Retries field	The number of times an I/O request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255.
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, update the following fields:

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

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

Name	Description
I/O Throttle Count field	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1,024.
LUNs per Target field	The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation. Enter an integer between 1 and 1,024. The recommended value is 1024.

Step 12 In the **Fibre Channel Port FLOGI** area, update the following fields:

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

Step 13 In the **Fibre Channel Port PLOGI** area, update the following fields:

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

Step 14 In the **SCSI I/O** area, update the following fields:

Name	Description
CDB Transmit Queue Count field	The number of SCSI I/O queue resources the system should allocate. Enter an integer between 1 and 8.
CDB Work Queue Ring Size 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
FC Work Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 128.
FC Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 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 12.

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

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
Index column	The unique identifier for the binding.
Target WWPN column	The target World Wide Port Name with which the binding is associated.
Host WWPN column	The host World Wide Port Name with which the binding is associated.
Bus ID column	The bus ID with which the binding is associated.
Target ID column	The target ID on the host system with which the binding is associated.
Rebuild Persistent Bindings button	Clears all unused bindings and resets the ones that are in use.
Close 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 VIC1225 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.

Viewing 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, 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.
MTU field	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9000.
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	The MAC address associated with the vNIC. To let the adapter select an available MAC address from its internal pool, select Auto . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
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. Note 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 Order field	The order in which this vNIC will be used. To let the system set the order, select Any . To specify an order, select the second radio button and enter an integer between 0 and 17.

Name	Description
Default VLAN field	<p>If there is no default VLAN for this vNIC, click NONE. Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p>Note This option cannot be used in VNTAG mode.</p>
VLAN Mode drop-down list	<p>If you want to use VLAN trunking, select TRUNK. Otherwise, select ACCESS.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Rate Limit 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.</p> <p>Note This option cannot be used in VNTAG mode.</p>
Enable PXE Boot check box	<p>Check this box if the vNIC can be used to perform a PXE boot.</p>
Channel Number field	<p>Select the channel number that will be assigned to this vNIC.</p> <p>Note VNTAG mode is required for this option.</p>
Port Profile 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>Note VNTAG mode is required for this option.</p>
Enable Uplink Failover 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>Note VNTAG mode is required for this option.</p>
Failback Timeout 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>Note VNTAG mode is required for this option.</p>

Step 9 In the **Ethernet Interrupt** area, review the information in the following fields:

Name	Description
Interrupt Count field	The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources. Enter an integer between 1 and 514.
Coalescing Time field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the Coalescing Time field.
Interrupt Mode drop-down list	The preferred driver interrupt mode. This can be one of the following: <ul style="list-style-type: none"> • MSI-X—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—MSI only. • INTx—PCI INTx interrupts.

Step 10 In the **Ethernet Receive Queue** area, review the information in the following fields:

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

Step 11 In the **Ethernet Transmit Queue** area, review 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.

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

Step 12 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 13 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. Note This option is also known as Large Send Offload (LSO).
Enable TCP Rx Offload Checksum Validation check box	If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. If cleared, the CPU validates all packet checksums.
Enable TCP Tx Offload Checksum Generation check box	If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. If cleared, the CPU calculates all packet checksums.
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 14 In the **Receive Side Scaling** area, review the information in the following fields:

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

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.
MTU field	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9000.
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	The MAC address associated with the vNIC. To let the adapter select an available MAC address from its internal pool, select Auto . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
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. Note 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 Order field	The order in which this vNIC will be used. To let the system set the order, select Any . To specify an order, select the second radio button and enter an integer between 0 and 17.
Default VLAN field	If there is no default VLAN for this vNIC, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field. Note This option cannot be used in VNTAG mode.
VLAN Mode drop-down list	If you want to use VLAN trunking, select TRUNK . Otherwise, select ACCESS . Note This option cannot be used in VNTAG mode.
Rate Limit field	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. Enter an integer between 1 and 10,000 Mbps. Note This option cannot be used in VNTAG mode.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.

Name	Description
Channel Number field	Select the channel number that will be assigned to this vNIC. Note VNTAG mode is required for this option.
Port Profile drop-down list	Select the port profile that should be associated with the vNIC. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.
Enable Uplink Failover check box	Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems. Note VNTAG mode is required for this option.
Failback Timeout field	After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC. Enter a number of seconds between 0 and 600. Note VNTAG mode is required for this option.

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

Name	Description
Interrupt Count field	The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources. Enter an integer between 1 and 514.
Coalescing Time field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the Coalescing Time field.

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

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

Name	Description
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
Receive Queue Ring Size 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
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Transmit Queue Ring Size 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
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 13 In the **TCP Offload** area, update the following fields:

Name	Description
Enable TCP Segmentation Offload check box	<p>If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate.</p> <p>If cleared, the CPU segments large packets.</p> <p>Note This option is also known as Large Send Offload (LSO).</p>
Enable TCP Rx Offload Checksum Validation 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>
Enable TCP Tx Offload Checksum Generation 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>
Enable Large Receive check box	<p>If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput.</p> <p>If cleared, the CPU processes all large packets.</p>

Step 14 In the **Receive Side Scaling** area, update the following fields:

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

Step 15 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 26.

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 adapter, such as the Cisco UCS VIC-1280, 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.

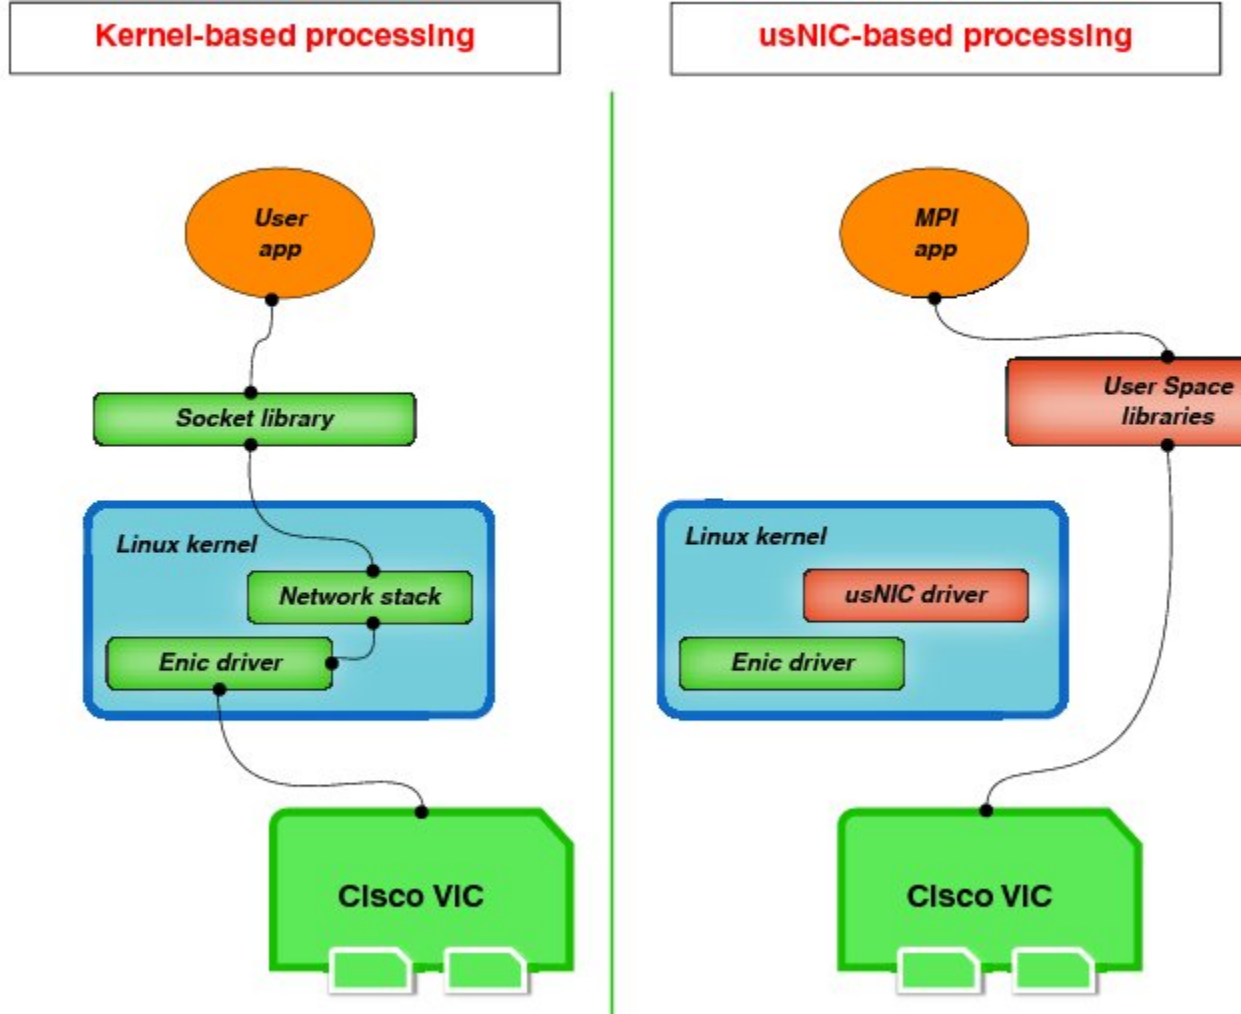
Cisco usNIC offers the following benefits for your MPI applications:

- 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 adapter, such as the Cisco UCS VIC-1280
 - 10GbE network

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 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.

Procedure

- Step 1** Log into the Cisco IMC GUI.
For more information about how to log into Cisco IMC, see the *Cisco UCS C-Series Servers Integrated Management Controller GUI Configuration Guide* available at this URL: http://www.cisco.com/en/US/products/ps10739/products_installation_and_configuration_guides_list.html.
- 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
Transmit Queue Count	The number of transmit queue resources to allocate. MPI will use 2 transmit queues per process. Therefore, Cisco recommends that you set this value to 2.
Receive Queue Count	The number of receive queue resources to allocate. MPI will use 2 receive queues per process. Therefore, Cisco recommends that you set this value to 2.
Completion Queue Count	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. Cisco recommends that you set this value to 4.

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. Note This field is read-only.
usNIC field	The number of usNICs assigned to the specific vNIC. Enter an integer between 0 and 225. To assign additional usNICs to a specified vNIC, enter value higher than the existing value. To delete usNICs from a specified vNIC, enter value smaller than the existing value. To delete all the usNICs assigned to a vNIC, enter zero.

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.

Name	Description
Interrupt Count field	The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources. Enter an integer between 1 and 514.
Interrupt Coalescing Type drop-down list	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as long as the time specified in the Coalescing Time field.
Interrupt Coalescing Timer Time field	The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.
Class of Service field	The class of service to associate with traffic from this usNIC. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. Note This option cannot be used in VNTAG mode.
TCP Segment Offload check box	If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. If cleared, the CPU segments large packets. Note This option is also known as Large Send Offload (LSO).
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.
TCP Tx Checksum 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.

Name	Description
TCP Rx Checksum 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.

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

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

Step 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.
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**.
-

Managing VM FEX

Virtual Machine Fabric Extender

Cisco Virtual Machine Fabric Extender (VM FEX) extends the (prestandard) IEEE 802.1Qbh port extender architecture to virtual machines. In this architecture, each VM interface is provided with a virtual Peripheral Component Interconnect Express (PCIe) device and a virtual port on a switch.

For this release, VM FEX supports the following cards and Operating systems:

Cards - Cisco UCS 1225 Virtual Interface Card

Operating Systems:

- VMware ESXi 5.1 Update 2
- VMware ESXi 5.5

VM FEX is not supported on Microsoft Hyper-V and Red Hat KVM for this release.

Viewing Virtual FEX 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, 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 **VM FEXs** tab.

Step 6 In the Virtual FEXs area, review the following information:

Name	Description
Properties button	Opens a dialog box that allows you to view the properties for the selected VM FEX.
Name column	The name of the VM FEX.
MTU column	The maximum transmission unit, or packet size, that this VM FEX accepts.
CoS column	If enabled, the VM FEX uses the class of service provided by the host operating system.
VLAN column	The VLAN associated with the VM FEX.
VLAN Mode column	The mode for the associated VLAN.
Uplink Failover column	If VNTAG mode is enabled for the adapter, this column displays whether traffic on this VM FEX will fail over to a secondary interface if the primary interface fails.

Step 7 In the Virtual FEXs area, select a VM FEX from the table.

Step 8 Click **Properties** to open the **VM FEX Properties** dialog box for the selected VM FEX.

Step 9 In the **General Properties** area, review the information in the following fields:

Name	Description
Name field	The name of the VM FEX.
MTU field	The maximum transmission unit, or packet size, that this VM FEX accepts.
Trust Host CoS field	If enabled, the VM FEX uses the class of service provided by the host operating system.
PCI Order field	The order in which this VM FEX will be used, if any.
Default VLAN field	The VLAN associated with the VM FEX.
Rate Limit field	The data rate limit associated with this VM FEX, if any.
PXE Boot field	Whether PXE boot is enabled or disabled for this VM FEX.

Step 10 In the **Ethernet Interrupt** area, review the information in the following fields:

Name	Description
Interrupt Count field	The number of interrupt resources allocated to this VM FEX.
Coalescing Time field	The time Cisco IMC waits between interrupts or the idle period that must be encountered before an interrupt is sent.
Coalescing Type field	This can be one of the following: <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the Coalescing Time field.
Interrupt Mode field	The preferred driver interrupt mode. This can be one of the following: <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. • MSI—MSI only. • INTx—PCI INTx interrupts.

Step 11 In the **Ethernet Receive Queue** area, review the information in the following fields:

Name	Description
Receive Queue Count field	The number of receive queue resources allocated to this VM FEX.

Name	Description
Receive Queue Ring Size field	The number of descriptors in each receive queue.

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

Name	Description
Transmit Queue Count field	The number of transmit queue resources allocated to this VM FEX.
Transmit Queue Ring Size field	The number of descriptors in each transmit queue.

Step 13 In the **Completion Queue** area, review the information in the following fields:

Name	Description
Completion Queue Count field	The number of completion queue resources allocated to this VM FEX.
Completion Queue Ring Size field	The number of descriptors in each completion queue.

Step 14 In the **TCP Offload** area, review the information in the following fields:

Name	Description
Enable TCP Segmentation Offload field	If enabled, the CPU sends large TCP packets to the hardware to be segmented. If disabled, the CPU segments large packets. Note This option is also known as Large Send Offload (LSO).
Enable TCP Rx Offload Checksum Validation field	If enabled, the CPU sends all packet checksums to the hardware for validation. If disabled, the CPU validates all packet checksums.
Enable TCP Tx Offload Checksum Generation field	If enabled, the CPU sends all packets to the hardware so that the checksum can be calculated. If disabled, the CPU calculates all packet checksums.
Enable Large Receive field	If enabled, the hardware reassembles all segmented packets before sending them to the CPU. If disabled, the CPU processes all large packets.

Step 15 In the **Receive Side Scaling** area, review the information in the following fields:

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

Managing Storage Adapters

Create Virtual Drive from Unused Physical Drives

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
- Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
- Step 4** In the **Actions** area, click **Create Virtual Drive from Unused Physical Drives**.
The **Create Virtual Drive from Unused Physical Drives** dialog box displays.
- Step 5** In the **Create Virtual Drive from Unused Physical Drives** dialog box, select the RAID level for the new virtual drives:
This can be one of the following:
 - **Raid 0**—Simple striping.

- **Raid 1**—Simple mirroring.
- **Raid 5**—Striping with parity.
- **Raid 6**—Striping with two parity drives.
- **Raid 10**—Spanned mirroring.
- **Raid 50**—Spanned striping with parity.
- **Raid 60**—Spanned striping with two parity drives.

Step 6 In the **Create Drive Groups** area, choose one or more physical drives to include in the group. Use the >> button to add the drives to the **Drive Groups** table. Use the << button to remove physical drives from the drive group.

Note The size of the smallest physical drive in the drive group defines the maximum size used for all the physical drives. To ensure maximum use of space for all physical drives, it is recommended that the size of all the drives in the drive group are similar.

Step 7 In the **Virtual Drive Properties** area, update the following properties:

Name	Description
Name field	The name of the new virtual drive you want to create.
Read Policy drop-down list	The read-ahead cache mode. This value cannot be changed.
Cache Policy drop-down list	The cache policy used for buffering reads. This value cannot be changed.
Strip Size drop-down list	The size of each strip, in KB. This value cannot be changed.
Write Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Write Through— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to Write Through caching when the BBU cannot guarantee the safety of the cache in the event of a power failure. • Write Back Bad BBU—With this policy, write caching remains Write Back even if the battery backup unit is defective or discharged.

Name	Description
Size field	The size of the virtual drive you want to create. Enter a value and select one of the following units: <ul style="list-style-type: none"> • MB • GB • TB

Step 8 Click **Create Virtual Drive**.

Create Virtual Drive from an Existing Drive Group

Before You Begin

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

Procedure

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

Step 2 On the **Storage** tab, click the appropriate LSI MegaRAID controller.

Step 3 On the **Work** pane, click the **Controller Info** tab.

Step 4 In the **Actions** area, click **Create Virtual Drive from an Existing Virtual Drive Group**. The **Create Virtual Drive from an Existing Virtual Drive Group** dialog box displays.

Step 5 In the **Create Virtual Drive from an Existing Virtual Drive Group** dialog box, select the virtual drive whose drive group you want to use to create a new virtual drive.

Step 6 In the **Virtual Drive Properties** area, update the following properties:

Name	Description
Name field	The name of the new virtual drive you want to create.
Read Policy drop-down list	The read-ahead cache mode. This value cannot be changed.
Cache Policy drop-down list	The cache policy used for buffering reads. This value cannot be changed.
Strip Size drop-down list	The size of each strip, in KB. This value cannot be changed.

Name	Description
Write Policy drop-down list	<p>This can be one of the following</p> <ul style="list-style-type: none"> • Write Through— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to Write Through caching when the BBU cannot guarantee the safety of the cache in the event of a power failure. • Write Back Bad BBU—With this policy, write caching remains Write Back even if the battery backup unit is defective or discharged.
Size field	<p>The size of the virtual drive you want to create. Enter a value and select one of the following units:</p> <ul style="list-style-type: none"> • MB • GB • TB

Step 7 Click **Create Virtual Drive**.

Importing Foreign Configuration

When one or more physical drives that have previously been configured with a different controller are inserted into a server, they are identified as foreign configurations. You can import these foreign configurations to a controller.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Import Foreign Config**.
 - Step 5** Click **OK** to confirm.
-

Clearing Foreign Configuration



Important

This task clears all foreign configuration on the controller. Also, all configuration information from all physical drives hosting foreign configuration is deleted. This action cannot be reverted.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Clear Foreign Config**.
 - Step 5** Click **OK** to confirm.
-

Clearing a Boot Drive



Important

This task clears the boot drive configuration on the controller. This action cannot be reverted.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Clear Boot Drive**.
 - Step 5** Click **OK** to confirm.
-

Enabling JBOD

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage Adapters** pane, click the appropriate **MegaRAID** controller.
 - Step 3** On the **Work** pane, click **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Enable JBOD**.
 - Step 5** Click **Ok** to confirm.
-

Disabling JBOD

Before You Begin

JBOD option must be enabled for the selected controller.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage Adapters** pane, click the appropriate **MegaRAID** controller.
 - Step 3** On the **Work** pane, click **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Disable JBOD**.
 - Step 5** Click **Ok** to confirm.
-

Preparing a Drive for Removal



Note You can perform this task only on physical drives that display the **Unconfigured Good** status.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the drive you want to remove.
 - Step 5** In the **Actions** area, click **Prepare for Removal**.
 - Step 6** Click **OK** to confirm.
-

Retrieving TTY Logs for a Controller

This task retrieves the TTY logs for the controller and places it in the `/var/log` location. This ensures that this log data is available when Technical Support Data is requested.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Controller Info** tab.
 - Step 4** In the **Actions** area, click **Get TTY Log**.
 - Step 5** Click **OK** to confirm.
- Important** Retrieving TTY logs for a controller could take up to 2-4 minutes. Until this process is complete, do not initiate exporting technical support data.
-

Undo Preparing a Drive for Removal

Before You Begin

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

Procedure

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select a drive with a status of **Ready to Remove**.
 - Step 5** In the **Actions** area, click **Undo Prepare for Removal**.
 - Step 6** Click **OK** to confirm.
-

Making a Dedicated Hot Spare

Before You Begin

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

Procedure

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the physical drive you want to make a dedicated hot spare.
 - Step 5** In the **Actions** area, click **Make Dedicated Hot Spare**.
The **Make Dedicated Hot Spare** dialog box displays.
 - Step 6** In the **Virtual Drive Details** area, update the following properties:

Name	Description
Virtual Drive Number drop-down list	Select the virtual drive to which you want to dedicate the physical drive as hot spare.
Virtual Drive Name field	The name of the selected virtual drive.
Physical Drive Number field	The number of the physical drive.

Step 7 Click **Make Dedicated Hot Spare** to confirm.

Making a Global Hot Spare

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the physical drive you want to make a global hot spare.
 - Step 5** In the **Actions** area, click **Make Global Hot Spare**.
-

Removing a Drive from Hot Spare Pools

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the global or dedicated hot spare you want to remove from the hot spare pools.
 - Step 5** In the **Actions** area, click **Remove From Hot Spare Pools**.
-

Toggling Physical Drive Status

Before You Begin

- You must log in with admin privileges to perform this task.
- The controller must support the JBOD mode and the JBOD mode must be enabled.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the drive you want to set as unconfigured good.
 - Step 5** In the **Actions** area, click **Set State as Unconfigured Good**.
 - Step 6** Click **OK** to confirm that the JBOD mode be disabled.
The **Set State as JBOD** option is enabled.
 - Step 7** To enable the JBOD mode for the physical drive, click **Set State as JBOD**.
 - Step 8** Click **OK** to confirm.
The **Set State as Unconfigured Good** option is enabled.
-

Setting a Physical Drive as a Controller Boot Drive

Before You Begin

- You must log in with admin privileges to perform this task.
- The controller must support the JBOD mode and the JBOD mode must be enabled.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the drive you want to set as boot drive for the controller.
 - Step 5** In the **Actions** area, click **Set as Boot Drive**.
 - Step 6** Click **OK** to confirm.
-

Initializing a Virtual Drive

All data on a virtual drive is lost when you initialize the drive. Before you run an initialization, back up any data on the virtual drive that you want to save.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
- Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, choose the drive that you want to initialize.
- Step 5** In the **Actions** area, click **Initialize**.
The **Initialize Virtual Drive** dialog box displays.
- Step 6** Choose the type of initialization you want to use for the virtual drive.
This can be one of the following:
- **Fast Initialize**—This option allows you to start writing data to the virtual drive immediately.
 - **Full Initialize**—A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete.
- Step 7** Click **Initialize VD** to initialize the drive, or **Cancel** to close the dialog box without making any changes.
- Step 8** To view the status of the task running on the drive, in the **Operations** area, click **Refresh**.
The following details are displayed:

Name	Description
Operation	Name of the operation that is in progress on the drive.
Progress in %	Progress of the operation, in percentage complete.
Elapsed Time in secs	The number of seconds that have elapsed since the operation began.

Set as Boot Drive

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
 - Step 4** In the **Virtual Drives** area, choose the drive from which the controller must boot.
 - Step 5** In the **Actions** area, click **Set as Boot Drive**.
 - Step 6** Click **OK** to confirm.
-

Editing a Virtual Drive

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage Adapters** pane, click **LSI MegaRAID SAS 9266-8i**.
- Step 3** On the **Work** pane, click **Virtual Drive Info** tab.
- Step 4** In the **Actions** area, click **Edit Virtual Drive**.
- Step 5** Review the instructions, and then click **OK**.
The Edit Virtual Drive dialog box displays.
- Step 6** From the **Select RAID Level to migrate** drop-down list, choose a RAID level.
See the following table for RAID migration criteria:

Name	Description
<p>Select RAID Level to migrate drop-down list</p>	<p>Select the RAID level to which you want to migrate. Migrations are allowed for the following RAID levels:</p> <ul style="list-style-type: none"> • RAID 0 to RAID 1 • RAID 0 to RAID 5 • RAID 0 to RAID 6 • RAID 1 to RAID 0 • RAID 1 to RAID 5 • RAID 1 to RAID 6 • RAID 5 to RAID 0 • RAID 6 to RAID 0 • RAID 6 to RAID 5 <p>When you are migrating from one raid level to another, the data arms of the new RAID level should be equal to or greater than the existing one.</p> <p>In case of RAID 6, the data arms will be number of drives minus two, as RAID 6 has double distributed parity. For example, when you create RAID 6 with eight drives, the number of data arms will be $8 - 2 = 6$. In this case, if you are migrating from RAID 6 to RAID 0, RAID 0 must have a minimum of six drives. If you select lesser number of drives then Edit or Save button will be disabled.</p> <p>If you are adding, you can migrate to RAID 0 as you will not be deleting any drives.</p> <p>Note RAID level migration is not supported in the following cases:</p> <ul style="list-style-type: none"> • When there are multiple virtual drives in a RAID group. • With a combination of SSD/HDD RAID groups.

Step 7 From the **Write Policy** drop-down list in the **Virtual Drive Properties** area, choose one of the following:

- **Write Through**— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache.
- **Write Back**— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to **Write Through** caching when the BBU cannot guarantee the safety of the cache in the event of a power failure.
- **Write Back Bad BBU**—With this policy, write caching remains **Write Back** even if the battery backup unit is defective or discharged.

Step 8 Click **Save Changes**.

Deleting a Virtual Drive



Important This task deletes a virtual drive, including the drives that run the booted operating system. So back up any data that you want to retain before you delete a virtual drive.

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
 - Step 4** In the **Virtual Drives** area, select the virtual drive you want to delete.
 - Step 5** In the **Actions** area, click **Delete Virtual Drive**.
 - Step 6** Click **OK** to confirm.
-

Enabling Auto Learn Cycle for a Battery Backup Unit

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Battery Backup Unit** tab.
 - Step 4** From the **Actions** pane, click **Enable Auto Learn Mode**.
A dialog prompts you to confirm the task.
 - Step 5** Click **OK**.
-

Disabling Auto Learn Cycle for a Battery Backup Unit

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Battery Backup Unit** tab.
 - Step 4** From the **Actions** pane, click **Disable Auto Learn Mode**.
A dialog prompts you to confirm the task.
 - Step 5** Click **OK**.
-

Starting Learn Cycles for a Battery Backup Unit

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Battery Backup Unit** tab.
 - Step 4** From the **Actions** pane, click **Start Learn Cycle**.
A dialog prompts you to confirm the task.
 - Step 5** Click **OK**.
-

Toggling Locator LED for a Physical Drive

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
 - Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
 - Step 3** On the **Work** pane, click the **Physical Drive Info** tab.
 - Step 4** From the **Status** area, select **Turn On** or **Turn Off** radio button for the **Locator LED** field.
-

Viewing Storage Controller Logs

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click the appropriate LSI MegaRAID controller.
- Step 3** On the **Work** pane, click **Storage Log** tab and review the following information:

Name	Description
Time column	The date and time the event occurred.
Severity column	The event severity. This can be one of the following: <ul style="list-style-type: none"> • Emergency • Alert • Critical • Error • Warning • Notice • Informational • Debug
Description column	A description of the event.

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
Export to drop-down list	The remote server type. This can be one of the following: <ul style="list-style-type: none"> • TFTP Server • FTP Server • SFTP Server • SCP Server • HTTP Server
Server IP/Hostname field	The IP address or hostname of the server to which the adapter configuration file will be exported. Depending on the setting in the Export to drop-down list, the name of the field may vary.

Name	Description
Path and Filename field	The path and filename Cisco IMC should use when exporting the file to the remote server.
Username	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or HTTP.
Password	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
Import from drop-down list	The remote server type. This can be one of the following: <ul style="list-style-type: none"> • TFTP Server • FTP Server • SFTP Server • SCP Server • HTTP Server

Name	Description
Server IP/Hostname field	The IP address or hostname of the server on which the adapter configuration file resides. Depending on the setting in the Import from drop-down list, the name of the field may vary.
Path and Filename field	The path and filename of the configuration file on the remote server.
Username	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or HTTP.
Password	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 GUI or CLI. 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
Install from drop-down list	The remote server type. This can be one of the following: <ul style="list-style-type: none"> • TFTP Server • FTP Server • SFTP Server • SCP Server • HTTP Server
Server IP/Hostname field	The IP address or hostname of the server on which the adapter configuration file resides. Depending on the setting in the Install from drop-down list, the name of the field may vary.
Path and Filename field	The path and filename of the configuration file on the remote server.
Username	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or HTTP.
Password	The password for the remote server username. This field does not apply if the protocol is TFTP or HTTP.
Back button	Click this button if you want to specify a local path for the firmware package.
Install Firmware button	Click this button to install the selected firmware package in the adapter's backup memory slot.

Name	Description
Close 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.
-