Managing Network Adapters

This chapter includes the following sections:

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 polices 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.
- 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 **Network 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.
CIMC Management Enabled column	Whether the adapter is able to manage CIMC. 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.
CIMC Management Enabled field	If this field displays **yes**, then the adapter is functioning in Cisco Card Mode and passing CIMC management traffic through to the server CIMC.

Name	Description
Configuration Pending field	If this field displays **yes**, the adapter configuration has changed in CIMC 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.
NIV Mode field	Whether Network Interface Virtualization (NIV) is enabled. If NIV mode is enabled: • 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

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: • **Fault** • **Link Up** • **Link Down** • **SFP ID Error** • **SFP Not Installed** • **SFP Security Check Failed** • **Unsupported SFP**
Encap column	The attribute added to the virtual network tag (VNTag) to support Network Interface Virtualization (NIV).

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

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 21
- Viewing Virtual FEX Properties, on page 31
- Viewing vHBA Properties, on page 7

Configuring Adapter Properties

Before You Begin

- You must log in with admin privileges to perform this task.
- 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 **Network 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 **Set Adapter Properties**.
The **Modify Adapter Properties** dialog box opens.

Step 7 In the **Modify Adapter Properties** dialog box, update the following fields:

Name	Description
Description field	A user-defined description for the adapter. You can enter between 1 and 63 characters.
Enable FIP Mode check box	If checked, then FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards. **Note** We recommend that you use this option only when explicitly directed to do so by a technical support representative.
Enable NIV Mode check box	If checked, then Network Interface Virtualization (NIV) mode is enabled. If NIV mode is enabled: • 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
Number of VM FEX Interfaces field	The number of VM FEX interfaces you want CIMC to create. Enter an integer between 0 and 112. **Note** NIV mode is required for this option.

Step 8 Click **Save Changes**.

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 11 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 **Network 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.
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.

Name	Description
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 NIV 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 NIV 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.

Name	Description
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** NIV 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** NIV 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.
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: • **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.

Name	Description
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 **Network 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 NIV 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 NIV 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** NIV 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** NIV 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: • **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 **Network 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 11.

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 **Network 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 **Network 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 **Network 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 **Network 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 **Network 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 **Network 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 NIV 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	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 NIV 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 NIV 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 NIV mode.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Channel Number field	Select the channel number that will be assigned to this vNIC. **Note** NIV 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** NIV 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** NIV 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** NIV mode is required for this option.

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: • **MIN**—The system waits for the time specified in the **Coalescing Time** field before sending another interrupt event. • **IDLE**—The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the **Coalescing Time** field.
Interrupt Mode drop-down list	The preferred driver interrupt mode. This can be one of the following: • **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 **Network 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 NIV 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 NIV 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 NIV 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 NIV 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** NIV 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** NIV 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** NIV 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** NIV 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: • **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: • **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	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, update the following fields:

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

Step 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 **Network 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 25.

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 **Network 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 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.

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 **Network 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 NIV 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 CIMC 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: • **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: • **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.
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.

Name	Description
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.

Backing Up and Restoring the Adapter Configuration

Exporting the Adapter Configuration

The adapter configuration can be exported as an XML file to a TFTP server.

Before You Begin

Obtain the TFTP 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 **Network 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
TFTP Server IP Address field	The IP address of the TFTP server to which the adapter configuration file will be exported.
Path and Filename field	The path and filename CIMC should use when exporting the file to the TFTP server.

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 **Network 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
TFTP Server IP Address field	The IP address of the TFTP server on which the adapter configuration file resides.
Path and Filename field	The path and filename of the configuration file on the TFTP server.

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 **Network 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 CIMC 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 CIMC 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 **Network 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 TFTP 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 **Network 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 TFTP server**, then click **Next**.

Step 8 In the **Install Adapter Firmware** dialog box, update the following fields:

Name	Description
TFTP Server IP Address field	The IP address of the TFTP server on which the adapter configuration file resides.
Path and Filename field	The path and filename of the configuration file on the TFTP server.
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.
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 **Network 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 **Network 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.