



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](#)
- [Configuring Adapter Properties, page 4](#)
- [Managing vHBAs, page 5](#)
- [Managing vNICs, page 15](#)
- [Backing Up and Restoring the Adapter Configuration, page 25](#)
- [Managing Adapter Firmware, page 27](#)

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. Following are the available adapters:

- Cisco UCS P81E Virtual Interface Card

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 2 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.

Viewing Network Adapter Properties

Before You Begin

- The server must be powered on, or the properties will not display.
- A Cisco UCS P81E Virtual Interface Card 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.
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: <ul style="list-style-type: none"> • Can be assigned to a specific channel • Can be associated with a port profile • Can fail over to another vNIC if there are communication problems

Step 8 In the **Uplinks** 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"> • Link Up

Name	Description
	<ul style="list-style-type: none"> • Link Down • Unsupported Transceiver
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	<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.
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 and virtual HBAs, see *Viewing vNIC Properties* and *Viewing vHBA Properties*.

Configuring Adapter Properties

Before You Begin

- You must log in with admin privileges to perform this task.
- A Cisco UCS P81E Virtual Interface Card 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: <ul style="list-style-type: none"> • Can be assigned to a specific channel • Can be associated with a port profile • Can fail over to another vNIC if there are communication problems

- 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 provides two vHBAs (fc0 and fc1). You cannot create additional vHBAs on this adapter card.

- When using the Cisco UCS P81E Virtual Interface Card in an FCoE application, you must associate the vHBA with the FCoE VLAN. Follow the instructions in to assign the VLAN.
- You must reset the adapter card after making configuration changes.

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 **Virtual HBAs** 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 system-assigned name of the virtual HBA.
World Wide Node Name field	The WWNN associated with the vHBA.
World Wide Port Name field	The WWPN associated with the vHBA.
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 field	The uplink port associated with the vHBA.
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 in the field. The ID can be an integer between 1 and 4094.
Class of Service drop-down list	Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.

Name	Description
Rate Limit field	If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter a rate limit in the associated field. You can enter an integer between 1 and 10,000 Mbps.
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 99.
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.

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: <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 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 1024.
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 1024. 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 **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 **Virtual HBAs** 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 system-assigned name of the virtual HBA.
World Wide Node Name field	The WWNN associated with the vHBA.
World Wide Port Name field	The WWPN associated with the vHBA.

Name	Description
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 field	The uplink port associated with the vHBA.
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 in the field. The ID can be an integer between 1 and 4094.
Class of Service drop-down list	Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.
Rate Limit field	If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter a rate limit in the associated field. You can enter an integer between 1 and 10,000 Mbps.
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 99.
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.

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

Name	Description
Enable FCP Error Recovery check box	If checked, the system uses FCP Sequence Level Error Recovery protocol (FC-TAPE).
Link Down Timeout field	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"> • MSI-X—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 1024.
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 1024. 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.

Name	Description
	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**.

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 **Virtual HBAs** 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 **Virtual HBAs** 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 **Virtual HBAs** 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.

Name	Description
Target WWPN column	The target World Wide Port Name with which the binding is associated.
Target LUN column	The target LUN ID with which the binding is associated.
Host LUN column	The target LUN ID on the host system with which the binding is associated.
Clear Persistent Bindings button	Clears all current bindings.
Close button	Closes the dialog box and saves your changes.

Step 9 Click Close.

Clearing 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 **Virtual HBAs** 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 **Clear 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 provides two default vNICs (eth0 and eth1). You can create up to 16 additional vNICs on this adapter card.

- You must reset the adapter card after making configuration changes.

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 **Virtual Ethernet Interface Cards** 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	A user-defined name for the virtual NIC. Once you create the vNIC, this name cannot be changed.
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.
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 99.

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 in the field. The ID can be an integer between 1 and 4094.
VLAN Mode drop-down list	If you want to use VLAN trunking, select TRUNK . Otherwise, select ACCESS .
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. You can enter an integer between 1 and 10,000 Mbps.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Channel Number field	If NIV mode is enabled for the adapter, select the channel number that will be assigned to this vNIC.
Port Profile drop-down list	If NIV mode is enabled for the adapter, select the port profile that should be associated with the vNIC. Note This field displays the port profiles defined on the switch to which this server is connected.
Enable Uplink Failover check box	If NIV mode is enabled for the adapter, check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems.
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.

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 a value 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.

Name	Description
	<ul style="list-style-type: none"> • 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	<p>The preferred driver interrupt mode. This can be one of the following:</p> <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	<p>The number of receive queue resources to allocate.</p> <p>Enter an integer between 1 and 256.</p>
Receive Queue Ring Size field	<p>The number of descriptors in each receive queue.</p> <p>Enter an integer between 64 and 4096.</p>

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

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

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

Name	Description
Completion Queue Count field	<p>The number of completion queue resources to allocate. In general, the number of completion queue resources you should allocate is equal to the number of transmit queue resources plus the number of receive queue resources.</p> <p>Enter an integer between 1 and 512.</p>
Completion Queue Ring Size field	<p>The number of descriptors in each completion queue.</p> <p>This value cannot be changed.</p>

Step 13 In the **TCP Offload** area, review the information in 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, review the information in the following fields:

Name	Description
Enable TCP Receive Side Scaling check box	<p>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.

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 **Virtual Ethernet Interface Cards** 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	A user-defined name for the virtual NIC. Once you create the vNIC, this name cannot be changed.
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.
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 99.

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 in the field. The ID can be an integer between 1 and 4094.
VLAN Mode drop-down list	If you want to use VLAN trunking, select TRUNK . Otherwise, select ACCESS .
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. You can enter an integer between 1 and 10,000 Mbps.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Channel Number field	If NIV mode is enabled for the adapter, select the channel number that will be assigned to this vNIC.
Port Profile drop-down list	If NIV mode is enabled for the adapter, select the port profile that should be associated with the vNIC. Note This field displays the port profiles defined on the switch to which this server is connected.
Enable Uplink Failover check box	If NIV mode is enabled for the adapter, check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems.
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.

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 a value 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.

Name	Description
	<ul style="list-style-type: none"> • 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	<p>The preferred driver interrupt mode. This can be one of the following:</p> <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	<p>The number of receive queue resources to allocate.</p> <p>Enter an integer between 1 and 256.</p>
Receive Queue Ring Size field	<p>The number of descriptors in each receive queue.</p> <p>Enter an integer between 64 and 4096.</p>

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

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

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

Name	Description
Completion Queue Count field	<p>The number of completion queue resources to allocate. In general, the number of completion queue resources you should allocate is equal to the number of transmit queue resources plus the number of receive queue resources.</p> <p>Enter an integer between 1 and 512.</p>
Completion Queue Ring Size field	<p>The number of descriptors in each completion queue.</p> <p>This value cannot be changed.</p>

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>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 **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 **Virtual Ethernet Interface Cards** 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*.

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 **Virtual Ethernet Interface Cards** area, select a vNIC from the table.
 - Step 7** In the **Virtual Ethernet Interface Cards** area, select a vNIC from the table.
Note You cannot delete either of the two default vNICs, **eth0** or **eth1**.
 - Step 8** Click **Delete** and click **OK** to confirm.
-

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

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**.
-

