



Managing Network Adapters

This chapter includes the following sections:

- [Viewing Network Adapter Properties, page 1](#)
- [Viewing Storage Adapter Properties, page 7](#)
- [Managing vHBAs, page 14](#)
- [Managing vNICs, page 27](#)
- [Backing Up and Restoring the Adapter Configuration, page 50](#)
- [Resetting the Adapter, page 53](#)

Viewing Network Adapter Properties

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Networking** menu, click **Adapter Card 1** or **Adapter Card 2**.
- Step 3** In the **Adapter Card Properties** area, review the following information:

Name	Description
PCI Slot field	The PCI slot in which the adapter is installed.
Vendor field	The vendor for the adapter.
Product Name field	The product name for the adapter.
Product ID field	The product ID for the adapter.

Name	Description
Serial Number field	The serial number for the adapter.
Version ID field	The version ID for the adapter.
Hardware Revision field	The hardware revision for the adapter.
Cisco IMC Management Enabled field	If this field displays yes , then the adapter is functioning in Cisco Card Mode and passing Cisco IMC management traffic through to the server Cisco IMC.
Configuration Pending field	If this field displays yes , the adapter configuration has changed in Cisco IMC but these changes have not been communicated to the host operating system. To activate the changes, an administrator must reboot the adapter.
iSCSI Boot Capable field	Whether iSCSI boot is supported on the adapter.
CDN Capable field	Whether CDN is supported on the adapter.
usNIC Capable field	Whether the adapter and the firmware running on the adapter support the usNIC.
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 LLDP check box	If checked, then Link Layer Discovery Protocol (LLDP) enables all the Data Center Bridging Capability Exchange protocol (DCBX) functionality, which includes FCoE, priority based flow control. By default, LLDP option is enabled. Note We recommend that you do not disable LLDP option, as it disables all the DCBX functionality. Note This option is available only on some UCS C-Series servers.
Enable VNTAG Mode check box	If VNTAG mode is enabled: <ul style="list-style-type: none"> • vNICs and vHBAs can be assigned to a specific channel. • vNICs and vHBAs can be associated to a port profile. • vNICs can fail over to another vNIC if there are communication problems.

Name	Description
<p>Port-0 drop-down list</p>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • 10 Gbps • 40 Gbps • 4 x 10 Gbps • Auto • 40 Gbps • 4 x 10 Gbps • Auto— For Cisco UCS VIC 13xx (or later generation) adapter series, the Auto value allows the adapter to determine the speed of the port based on the transceiver module inserted into the port. <p>Note On adapters that support 40 Gbps speed, Auto is the default option.</p> <p>Note You need to choose 40 Gbps as the port speed if you are using a 40 Gbps switch.</p>
<p>Port-1 drop-down list</p>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • 10 Gbps • 40 Gbps • 4 x 10 Gbps • Auto • 40 Gbps • 4 x 10 Gbps • Auto— For Cisco UCS VIC 13xx (or later generation) adapter series, the Auto value allows the adapter to determine the speed of the port based on the transceiver module inserted into the port. <p>Note On adapters that support 40 Gbps speed, Auto is the default option.</p> <p>Note You need to choose 40 Gbps as the port speed if you are using a 40 Gbps switch.</p>
<p>Training Link - 0 check box</p>	<p>If checked, link training for port-0 is enabled.</p> <p>Note Is supported only if port speed of 40 Gbps on port-0 is selected.</p>
<p>Training Link - 1 check box</p>	<p>If checked, link training for port-1 is enabled.</p> <p>Note Is supported only if port speed of 40 Gbps on port-1 is selected.</p>

Name	Description
PCI Link field	In a single server dual SIOC setup, depending on the information in the PCI slot field, this indicates the vNICs and vHBA of that adapter associated with the server. For example, if the PCI Slot displays 2, and the PCI Link field displays Server-1 as the field value, it indicates that vNICs and vHBA of adapter-2 are PCI linked to the server 1 host.

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

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.

Step 5 In the **External Ethernet Interfaces** area, review the following information:

Name	Description
ID column	The uplink port ID.
MAC Address column	The MAC address of the uplink port.

Name	Description
Link State column	<p>The current operational state of the uplink port. This can be one of the following:</p> <ul style="list-style-type: none"> • Fault • Link Up • Link Down • SFP ID Error • SFP Not Installed • SFP Security Check Failed • Unsupported SFP
Encap column	<p>The mode in which adapter operates. This can be one of the following:</p> <ul style="list-style-type: none"> • CE—Classical Ethernet mode. • NIV—Network Interface Virtualization mode.
Admin Speed column	<p>The data transfer rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> • 1 Gbps • 10 Gbps • 40 Gbps • 4 x 10 Gbps • Auto • 40 Gbps • 4 x 10 Gbps <p>Note You need to choose 40 Gbps as the port speed if you are using a 40 Gbps switch.</p>

Name	Description
Operating Speed column	<p>The operating rate for the port. This can be one of the following:</p> <ul style="list-style-type: none"> • 1 Gbps • 10 Gbps • 40 Gbps • 4 x 10 Gbps • Auto • 40 Gbps • 4 x 10 Gbps <p>Note You need to choose 40 Gbps as the port speed if you are using a 40 Gbps switch.</p>
Training Link column	Indicates if link training is enabled on the port.
Connector Present column	<p>Indicated whether or not the connector is present. This can be one of the following:</p> <ul style="list-style-type: none"> • Yes—Connector is present. • No—Connector not present. <p>Note This option is only available for some adapter cards.</p>
Connector Supported column	<p>Indicates whether or not the connector is supported by Cisco. This can be one of the following:</p> <ul style="list-style-type: none"> • Yes—The connector is supported by Cisco. • No—The connector is not supported by Cisco. <p>If the connector is not supported then the link will not be up.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Type column	<p>The type of the connector.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Vendor column	<p>The vendor for the connector.</p> <p>Note This option is only available for some adapter cards.</p>
Connector Part Number column	<p>The part number of the connector.</p> <p>Note This option is only available for some adapter cards.</p>

Name	Description
Connector Part Revision column	The part revision number of the connector. Note This option is only available for some adapter cards.

Viewing Storage Adapter Properties

Before You Begin

- The server must be powered on.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **Controller** area, the **Controller Info** tab displays by default.
- Step 4** In the **Work** pane's **Health/Status** area, review the following information:

Name	Description
Composite Health field	The combined health of the controller, the attached drives, and the battery backup unit. This can be one of the following: <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault • N/A
Controller Status field	The current status of the controller. This can be one of the following: <ul style="list-style-type: none"> • Optimal — The controller is functioning properly. • Failed — The controller is not functioning. • Unresponsive — The controller is down.
RAID Chip Temperature field	Temperature of the controller in degree centigrade.

Name	Description
TTY Log Status field	The current status of the TTY log download. This can be one of the following: <ul style="list-style-type: none"> • Not Downloaded • In Progress • Complete

Step 5 In the **Firmware Versions** area, review the following information:

Name	Description
Product Name field	The name of the MegaRAID controller.
Serial Number field	The serial number of the MegaRAID controller.
Firmware Package Build field	The active firmware package version number. For the firmware component version numbers, see the Running Firmware Images area.

Step 6 In the **PCI Info** area, review the following information:

Name	Description
PCI Slot field	The name of the PCIe slot in which the controller is located.
Vendor ID field	The PCI vendor ID, in hexadecimal.
Device ID field	The PCI device ID, in hexadecimal.
SubVendor ID field	The PCI subvendor ID, in hexadecimal.
SubDevice ID field	The PCI subdevice ID, in hexadecimal.

Step 7 In the **Manufacturing Data** area, review the following information:

Name	Description
Manufactured Date field	The date the MegaRAID card was manufactured, in the format yy-mm-dd.
Revision No field	The board revision number, if any.

Step 8 In the **Boot Drive** area, review the following information:

Name	Description
Boot Drive field	The number of the boot drive.
Boot Drive is PD field	If this field displays true , the boot drive is a physical drive.

Step 9 In the **Running Firmware Images** area, review the following information:

Name	Description
BIOS Version field	The BIOS option PROM version number.
Firmware Version field	The active firmware version number.
Preboot CLI Version field	The pre-boot CLI version number.
WebBIOS Version field	The Web BIOS version number.
NVDATA Version field	The non-volatile data (NVDATA) version number.
Boot Block Version field	The boot block version number.
Boot Version field	The firmware boot loader version number on the LSI controller.

Step 10 In the **Startup Firmware Images** area, review the following information:

Name	Description
Startup BIOS Version field	The BIOS option PROM version that will become active when the host server reboots, if different from the current version.
Startup Firmware Version field	The firmware version that will become active when the host server reboots, if different from the current version.
Startup Preboot CLI Version field	The pre-boot CLI version that will become active when the host server reboots, if different from the current version.
Startup WebBIOS Version field	The Web BIOS version that will become active when the host server reboots, if different from the current version.

Name	Description
Startup NVDATA Version field	The non-volatile data version that will become active when the host server reboots, if different from the current version.
Startup Boot Block Version field	The boot block version that will become active when the host server reboots, if different from the current version.
Startup Boot Version field	The firmware boot loader version that will become active when the host server reboots, if different from the current version.

Step 11 In the **Virtual Drive Count** area, review the following information:

Name	Description
Virtual Drive Count field	The number of virtual drives configured on the controller.
Degraded Drive Count field	The number of virtual drives in a degraded state on the controller.
Offline Drive Count field	The number of virtual drives that have failed on the controller.

Step 12 In the **Physical Drive Count** area, review the following information:

Name	Description
Disk Present Count field	The number of physical drives present on the controller.
Degraded Disk Count field	The number of physical drives in a degraded state on the controller.
Failed Disk Count field	The number of physical drives that have failed on the controller.

Step 13 In the **Settings** area, review the following information:

Name	Description
Predictive Fail Poll Interval field	<p>The number of seconds between predictive failure polls.</p> <p>During each poll, the controller examines the Self-Monitoring Analysis and Reporting Technology (SMART) data on all physical drives to determine if any is about to fail.</p>
Rebuild Rate field	<p>The rate at which the controller rebuilds degraded RAID volumes.</p> <p>This rate is shown as a percentage of the total bandwidth available.</p>
Patrol Read Rate field	<p>The rate at which the controller performs a background read of the physical drives looking for inconsistent data.</p> <p>This rate is shown as a percentage of the total bandwidth available.</p>
Consistency Check Rate field	<p>The rate at which the controller scans the virtual drives looking for redundant data inconsistencies and fixing them.</p> <p>This rate is shown as a percentage of the total bandwidth available.</p>
Reconstruction Rate field	<p>The rate at which virtual drives are reconstructed when the capacity or RAID level needs to be changed.</p> <p>This rate is shown as a percentage of the total bandwidth available.</p>
Cache Flush Interval field	<p>The number of seconds waits before flushing the cache memory to the physical drives.</p>
Max Drives To Spin Up At Once field	<p>The number of drives that can be spun up simultaneously after the server is powered on.</p>
Delay Among Spinup Groups field	<p>The number of seconds to wait before the controller spins up the next set of drives.</p>

Name	Description
Physical Drive Coercion Mode field	<p>Whether the controller rounds the size of physical drives down to a round number. This can be one of the following:</p> <ul style="list-style-type: none"> • None—The controller does not do any rounding. • 128 MB—Drive sizes are rounded down to the closest multiple of 128 MB. • 1GB—Drive sizes are rounded down to the closest multiple of 1GB.
Cluster Mode field	If this field displays true , the drives on this controller are shared with controllers on other servers.
Battery Warning field	If this field displays true , missing battery warnings are disabled.
ECC Bucket Leak Rate field	<p>The error correcting code (ECC) single-bit error bucket leak rate, in minutes.</p> <p>With ECC, the controller increments an error counter when it encounters a single bit error while reading from a physical drive. The controller decrements the error counter each time the number of minutes defined in this field passes.</p> <p>If the error counter reaches a system-defined maximum, the controller sends an event message to the system.</p>
Expose Enclosure Devices field	If this field displays true , enclosure devices are visible to the host drivers.
Maintain PD Fail History field	If this field displays true , the controller remembers which physical drives were determined to be bad across server reboots.
Enable Copyback on SMART field	If this field displays true , the controller copies the contents of the drive to a spare drive if Self-Monitoring Analysis and Reporting Technology (SMART) reports an error.
Enable Copyback to SSD on SMART Error field	If this field displays true , the controller copies the contents of an SSD card to a spare card if SMART reports an error.
Native Command Queuing field	If this field displays true , Native Command Queuing (NCQ) is disabled.

Name	Description
JBOD field	If this field displays true , JBOD is enabled.
Enable Spin Down of Unconfigured Drives field	If this field displays true , the controller spins down unconfigured drives.
Enable SSD Patrol Read field	If this field displays true , the controller performs patrol reads on SSD cards.
Auto Enhanced Import field	If this field displays true , foreign configurations are automatically imported when the controller boots.

Step 14 In the **Capabilities** area, review the following information:

Name	Description
RAID Levels Supported field	<p>The RAID levels supported by the controller. This can be one or more of the following:</p> <ul style="list-style-type: none"> • Raid 0—Simple striping. • Raid 1—Simple mirroring. • Raid 5—Striping with parity. • Raid 1E—Integrated offset strip mirroring • Raid 6—Striping with two parity drives. • Raid 10—Spanned mirroring. • Raid 50—Spanned striping with parity. • Raid 60—Spanned striping with two parity drives. • Raid srl-03—Spanned secondary RAID level • Raid 00—Spanned striping. • Raid 1e-rlq0—Integrated adjacent strip mirroring with no span. • Raid 1e0-rlq0—Integrated adjacent strip mirroring with span.

Step 15 In the **HW Configuration** area, review the following information:

Name	Description
SAS Address field	A MegaRAID controller can have up to 16 serial-attached SCSI (SAS) addresses. This field displays the first 8 SAS addresses, if they are in use.

Name	Description
BBU Present field	If this field displays true , the battery backup unit is present.
NVRAM Present field	If this field displays true , the NVRAM is present.
NVRAM Size field	The size of the NVRAM, in kilobytes.
Serial Debugger Present field	If this field displays true , a serial debugger is attached to the RAID card.
Memory Present field	If this field displays true , memory is present.
Flash Present field	If this field displays true , flash memory is present.
Flash Size field	The size of the flash memory, in megabytes.
Memory Size field	The size of the memory, in megabytes.
Cache Memory Size field	The size of the cache memory, in megabytes.
Number of Backend Ports field	The number of SATA or SAS ports on the controller.

Step 16 In the **Error Counters** area, review the following information:

Name	Description
Memory Correctable Errors field	The number of correctable errors in the controller memory.
Memory Uncorrectable Errors field	The number of uncorrectable errors in the controller memory.

Managing vHBAs

Guidelines for Managing vHBAs

When managing vHBAs, consider the following guidelines and restrictions:

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



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

- When using the Cisco UCS P81E Virtual Interface Card or Cisco UCS VIC 1225 Virtual Interface Card in an FCoE application, you must associate the vHBA with the FCoE VLAN. Follow the instructions in the **Modifying vHBA Properties** section to assign the VLAN.
- After making configuration changes, you must reboot the host for settings to take effect.

Viewing vHBA Properties

Procedure

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

Name	Description
Name field	The name of the virtual HBA. This name cannot be changed after the vHBA has been created.
Target WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
Target WWPNN field	The WWPNN associated with the vHBA. To let the system generate the WWPNN, select AUTO . To specify a WWPNN, click the second radio button and enter the WWPNN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
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.

Name	Description
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
Class of Service drop-down list	The CoS for the vHBA. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. Note This option cannot be used in VNTAG mode.
Rate Limit field	The data rate limit for traffic on this vHBA, in Mbps. If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter an integer between 1 and 10,000. Note This option cannot be used in VNTAG mode.
PCIe Device Order field	The order in which this vHBA will be used. To let the system set the order, select ANY . To specify an order, select the second radio button and enter an integer between 0 and 17.
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.

Name	Description
Channel Number field	The channel number that will be assigned to this vHBA. Enter an integer between 1 and 1,000. Note VNTAG mode is required for this option.
Port Profile drop-down list	The port profile that should be associated with the vHBA, if any. This field displays the port profiles defined on the switch to which this server is connected. Note VNTAG mode is required for this option.

Step 5 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.
I/O Timeout Retry field	The time period till which the system waits for timeout before retrying. When a disk does not respond for I/O within the defined timeout period, the driver aborts the pending command, and resends the same I/O after the timer expires. Enter an integer between 1 and 59.
Port Down Timeout field	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. Enter an integer between 0 and 240,000.

Step 6 In the **Fibre Channel Interrupt** area, review the information in the following fields:

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

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

Name	Description
I/O Throttle Count field	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1,024.
LUNs per Target field	The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation. Enter an integer between 1 and 1,024. The recommended value is 1024.
LUN Queue Depth field	The number of commands that the HBA can send or receive in a single chunk per LUN. This parameter adjusts the initial queue depth for all LUNs on the adapter. Default value is 20 for physical miniports and 250 for virtual miniports.

Step 8 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 9 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 10 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 Transmit Queue Ring Size field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

Step 11 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 **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 3** In the **vHBAs** pane, click **fc0** or **fc1**.
- Step 4** In the **General** area, update the following fields:

Name	Description
Name field	The name of the virtual HBA. This name cannot be changed after the vHBA has been created.
Target WWNN field	The WWNN associated with the vHBA. To let the system generate the WWNN, select AUTO . To specify a WWNN, click the second radio button and enter the WWNN in the corresponding field.
Target WWPN field	The WWPN associated with the vHBA. To let the system generate the WWPN, select AUTO . To specify a WWPN, click the second radio button and enter the WWPN in the corresponding field.
FC SAN Boot check box	If checked, the vHBA can be used to perform a SAN boot.
Enable Persistent LUN Binding check box	If checked, any LUN ID associations are retained in memory until they are manually cleared.
Uplink Port field	The uplink port associated with the vHBA. Note This value cannot be changed for the system-defined vHBAs fc0 and fc1.
MAC Address field	The MAC address associated with the vHBA. To let the system generate the MAC address, select AUTO . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Default VLAN field	If there is no default VLAN for this vHBA, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.
Class of Service drop-down list	The CoS for the vHBA. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. Note This option cannot be used in VNTAG mode.
Rate Limit field	The data rate limit for traffic on this vHBA, in Mbps. If you want this vHBA to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter an integer between 1 and 10,000. Note This option cannot be used in VNTAG mode.

Name	Description
PCIe Device Order field	<p>The order in which this vHBA will be used.</p> <p>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.</p>
EDTOV field	<p>The error detect timeout value (EDTOV), which is the number of milliseconds to wait before the system assumes that an error has occurred.</p> <p>Enter an integer between 1,000 and 100,000. The default is 2,000 milliseconds.</p>
RATOV field	<p>The resource allocation timeout value (RATOV), which is the number of milliseconds to wait before the system assumes that a resource cannot be properly allocated.</p> <p>Enter an integer between 5,000 and 100,000. The default is 10,000 milliseconds.</p>
Max Data Field Size field	<p>The maximum size of the Fibre Channel frame payload bytes that the vHBA supports.</p> <p>Enter an integer between 256 and 2112.</p>
Channel Number field	<p>The channel number that will be assigned to this vHBA.</p> <p>Enter an integer between 1 and 1,000.</p> <p>Note VNTAG mode is required for this option.</p>
Port Profile drop-down list	<p>The port profile that should be associated with the vHBA, if any.</p> <p>This field displays the port profiles defined on the switch to which this server is connected.</p> <p>Note VNTAG mode is required for this option.</p>

Step 5 In the **Error Recovery** area, update the following fields:

Name	Description
Enable FCP Error Recovery check box	<p>If checked, the system uses FCP Sequence Level Error Recovery protocol (FC-TAPE).</p>
Link Down Timeout field	<p>The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost.</p> <p>Enter an integer between 0 and 240,000.</p>

Name	Description
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.
I/O Timeout Retry field	The time period till which the system waits for timeout before retrying. When a disk does not respond for I/O within the defined timeout period, the driver aborts the pending command, and resends the same I/O after the timer expires. Enter an integer between 1 and 59.
Port Down Timeout field	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. Enter an integer between 0 and 240,000.

Step 6 In the **Fibre Channel Interrupt** area, update the following fields:

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

Step 7 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.
LUN Queue Depth field	The number of commands that the HBA can send or receive in a single chunk per LUN. This parameter adjusts the initial queue depth for all LUNs on the adapter. Default value is 20 for physical miniports and 250 for virtual miniports.

Step 8 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 9 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 10 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 Transmit Queue Ring Size field	The number of descriptors in each SCSI I/O queue. Enter an integer between 64 and 512.

Step 11 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 12 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 **Networking** menu.

Step 2 In the **Adapter Card** pane, click the **vHBAs** tab.

Step 3 In the **Host Fibre Channel Interfaces** area, choose one of these actions:

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

The **Add vHBA** dialog box appears.

Step 4 In the **Add vHBA** dialog box, enter a name for the vHBA in the **Name** entry box.

Step 5 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 19.

Deleting a vHBA

Procedure

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

Step 2 In the **Adapter Card** pane, click the **vHBAs** tab.

Step 3 In the **Host Fibre Channel Interfaces** area, select a vHBA or vHBAs from the table.

Note You cannot delete either of the two default vHBAs, **fc0** or **fc1**.

Step 4 Click **Delete vHBAs** 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 **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vHBAs** tab.
- Step 3** In the **Fibre Channel Interfaces** area, scroll down to the **Boot Table** area.
- Step 4** Click the **Add Boot Entry** button to open the **Add Boot Entry** dialog box.
- Step 5** In the **Add Boot Entry** dialog box, review the following information and perform the actions specified:

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.

Deleting a Boot Table Entry

Procedure

-
- Step 1** In the **Navigation** pane, click the **Networking** menu.
 - Step 2** In the **Adapter Card** pane, click the **vHBAs** tab.
 - Step 3** In the Fibre Channel Interfaces area, scroll down to the **Boot Table** area.
 - Step 4** In the **Boot Table** area, click the entry to be deleted.
 - Step 5** Click **Delete Boot Entry** and click **OK** to confirm.
-

vHBA Persistent Binding

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

Viewing Persistent Bindings

Procedure

-
- Step 1** In the **Navigation** pane, click the **Networking** menu.
 - Step 2** In the **Adapter Card** pane, click the **vHBAs** tab.
 - Step 3** In the **vHBAs** pane, click **fc0** or **fc1**.
 - Step 4** In the **Persistent Bindings** dialog box, 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 5 Click **Close**.

Rebuilding Persistent Bindings

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
 - Step 2** In the **Adapter Card** pane, click the **vHBAs** tab.
 - Step 3** In the **vHBAs** pane, click **fc0** or **fc1**.
 - Step 4** In the **Fibre Channel Interfaces** area, scroll down to the **Persistent Bindings** area.
 - Step 5** Click the **Rebuild Persistent Bindings** button.
 - Step 6** Click **OK** to confirm.
-

Managing vNICs

Guidelines for Managing vNICs

When managing vNICs, consider the following guidelines and restrictions:

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



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

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

Cisco C-series servers use Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) for packet transfers. RoCE defines the mechanism of performing RDMA over ethernet, based on the similar mechanism of RDMA over Infiniband. However, RoCE, with its performance oriented characteristics, delivers a superior performance compared to traditional network socket implementation because of the lower latency, lower CPU utilization and higher utilization of network bandwidth. RoCE meets the requirement of moving large amount of data across networks very efficiently.

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

- Queue Pairs
- Memory Regions

- Resource Groups

Guidelines and Limitations for SMB Direct with RoCE

- Microsoft SMB Direct with RoCE is supported:
 - On Windows 2012 R2.
 - On Windows 2016.
- Cisco UCS C-Series server does not support more than 4 RoCE-enabled vNICs per adapter.
- Cisco UCS C-Series server does not support RoCE with NVGRE, VXLAN, VMQ, or usNIC.
- Maximum number of queue pairs per adapter is 8192.
- Maximum number of memory regions per adapter is 524288.
- RoCE configuration is supported between Cisco adapters. Interoperability between Cisco adapters and third party adapters is not supported.



Important

It is required to configure the no-drop QOS policy settings at the switches in the RDMA traffic path.

Viewing vNIC Properties

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Ethernet Interfaces** pane's **vNIC Properties** area, review the information in the following fields:

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

Name	Description
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	The MAC address associated with the vNIC. To let the adapter select an available MAC address from its internal pool, select Auto . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Class of Service drop-down list	The class of service to associate with traffic from this vNIC. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. Note This option cannot be used in VNTAG mode.
Trust Host CoS check box	Check this box if you want the vNIC to use the class of service provided by the host operating system.
PCI Order field	The order in which this vNIC will be used. To specify an order, enter an integer within the displayed range.
Default VLAN field	If there is no default VLAN for this vNIC, click NONE . Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field. Note This option cannot be used in VNTAG mode.
VLAN Mode drop-down list	If you want to use VLAN trunking, select TRUNK . Otherwise, select ACCESS . Note This option cannot be used in VNTAG mode.
Rate Limit field	If you want this vNIC to have an unlimited data rate, select OFF . Otherwise, click the second radio button and enter a rate limit in the associated field. Enter an integer between 1 and 10,000 Mbps. You can enter an integer between 1 and 40,000 Mbps for VIC 13xx controllers. Note This option cannot be used in VNTAG mode.
Enable PXE Boot check box	Check this box if the vNIC can be used to perform a PXE boot.
Channel Number field	Select the channel number that will be assigned to this vNIC. Note VNTAG mode is required for this option.

Name	Description
PCI Link field	<p>The link through which vNICs can be connected. These are the following values:</p> <ul style="list-style-type: none"> • 0 - The first cross-edged link where the vNIC is placed. • 1 - The second cross-edged link where the vNIC is placed. <p>Note</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers.
Port Profile drop-down list	<p>Select the port profile that should be associated with the vNIC.</p> <p>This field displays the port profiles defined on the switch to which this server is connected.</p> <p>Note VNTAG mode is required for this option.</p>
Enable Uplink Failover check box	<p>Check this box if traffic on this vNIC should fail over to the secondary interface if there are communication problems.</p> <p>Note VNTAG mode is required for this option.</p>
Enable VMQ check box	<p>Check this box to enable Virtual Machine Queue (VMQ).</p> <p>Note Ensure that VMQ is not enabled when SR-IOV or netflow option is enabled on the adapter.</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>
Enable aRFS check box	<p>Check this box to enable Accelerated Receive Flow steering (aRFS).</p> <p>This option is available only on some Cisco UCS C-Series servers.</p>
Enable NVGRE check box	<p>Check this box to enable Network Virtualization using Generic Routing Encapsulation.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385 cards.
Enable VXLAN check box	<p>Check this box to enable Virtual Extensible LAN.</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385 cards.
Advanced Filter check box	<p>Check this box to enable advanced filter options in vNICs.</p>

Name	Description
Failback Timeout field	<p>After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC.</p> <p>Enter a number of seconds between 0 and 600.</p> <p>Note VNTAG mode is required for this option.</p>

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

Name	Description
Interrupt Count field	<p>The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources.</p> <p>Enter an integer between 1 and 514.</p>
Coalescing Time field	<p>The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent.</p> <p>Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.</p>
Coalescing Type drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as least 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 6 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>

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

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

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

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

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

Name	Description
Enable TCP Segmentation Offload check box	If checked, the CPU sends large TCP packets to the hardware to be segmented. This option may reduce CPU overhead and increase throughput rate. If cleared, the CPU segments large packets. 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.

Name	Description
Enable TCP Tx Offload Checksum Generation check box	If checked, the CPU sends all packets to the hardware so that the checksum can be calculated. This option may reduce CPU overhead. If cleared, the CPU calculates all packet checksums.
Enable Large Receive check box	If checked, the hardware reassembles all segmented packets before sending them to the CPU. This option may reduce CPU utilization and increase inbound throughput. If cleared, the CPU processes all large packets.

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

Name	Description
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 **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Ethernet Interfaces** pane's **vNIC Properties** area, update the following fields:

Name	Description
Name field	The name for the virtual NIC. This name cannot be changed after the vNIC has been created.
CDN field	The Consistent Device Name (CDN) that you can assign to the ethernet vNICs on the VIC cards. Assigning a specific CDN to a device helps in identifying it on the host OS. Note This feature works only when the CDN Support for VIC token is enabled in the BIOS.
MTU field	The maximum transmission unit, or packet size, that this vNIC accepts. Enter an integer between 1500 and 9000.
Uplink Port drop-down list	The uplink port associated with this vNIC. All traffic for this vNIC goes through this uplink port.
MAC Address field	The MAC address associated with the vNIC. To let the adapter select an available MAC address from its internal pool, select Auto . To specify an address, click the second radio button and enter the MAC address in the corresponding field.
Class of Service drop-down list	The class of service to associate with traffic from this vNIC. Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority. Note This option cannot be used in VNTAG mode.
Trust Host CoS check box	Check this box if you want the vNIC to use the class of service provided by the host operating system.

Name	Description
<p>PCI Link field</p>	<p>The link through which vNICs can be connected. These are the following values:</p> <ul style="list-style-type: none"> • 0 - The first cross-edged link where the vNIC is placed. • 1 - The second cross-edged link where the vNIC is placed. <p>Note</p> <ul style="list-style-type: none"> • This option is available only on some Cisco UCS C-Series servers. • This option is available only on C-Series servers with Cisco VIC 1385 cards.
<p>PCI Order field</p>	<p>The order in which this vNIC will be used.</p> <p>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.</p>
<p>Default VLAN field</p>	<p>If there is no default VLAN for this vNIC, click NONE. Otherwise, click the second radio button and enter a VLAN ID between 1 and 4094 in the field.</p> <p>Note This option cannot be used in VNTAG mode.</p>
<p>VLAN Mode drop-down list</p>	<p>If you want to use VLAN trunking, select TRUNK. Otherwise, select ACCESS. When the VLAN is set to ACCESS mode, any frame received from the specified default VLAN (1-4094) that is received from the switch with a TAG removes that TAG when it is sent to the host OS through the vNIC.</p> <p>Note This option cannot be used in VNTAG mode.</p>
<p>Rate Limit field</p>	<p>If you want this vNIC to have an unlimited data rate, select OFF. Otherwise, click the second radio button and enter a rate limit in the associated field.</p> <p>Enter an integer between 1 and 10,000 Mbps or 40,000 Mbps depending on the adapter card you choose.</p> <p>Note This option cannot be used in VNTAG mode.</p>
<p>Enable PXE Boot check box</p>	<p>Check this box if the vNIC can be used to perform a PXE boot.</p>
<p>Channel Number field</p>	<p>Select the channel number that will be assigned to this vNIC.</p> <p>Note VNTAG mode is required for this option.</p>

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

Step 5 In the **Ethernet Interrupt** area, update the following fields:

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

Step 6 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 7 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.

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

Step 8 In the **Completion Queue** area, 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 9 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 10 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 11 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 **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the **Host Ethernet Interfaces** area, choose one of these actions:
- To create a vNIC using default configuration settings, click **Add vNIC**.
 - To create a vNIC using the same configuration settings as an existing vNIC, select that vNIC and click **Clone vNIC**.

The **Add vNIC** dialog box appears.

- Step 4** In the **Add vNIC** dialog box, enter a name for the vNIC in the **Name** entry box.
- Step 5** (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 6** Click **Add vNIC**.

What to Do Next

If configuration changes are required, configure the new vNIC as described in [Modifying vNIC Properties, on page 34](#).

Deleting a vNIC

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the **Host Ethernet Interfaces** area, select a vNIC from the table.
- Note** You cannot delete either of the two default vNICs, **eth0** or **eth1**.
- Step 4** Click **Delete vNIC** and click **OK** to confirm.
-

Managing Cisco usNIC

Overview of Cisco usNIC

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


Cisco usNIC offers the following benefits for your MPI applications:

- Provides a low-latency and high-throughput communication transport.
- Employs the standard and application-independent Ethernet protocol.
- Takes advantage of lowlatency forwarding, Unified Fabric, and integrated management support in the following Cisco data center platforms:
 - Cisco UCS server

- Cisco UCS VIC second generation or later generation adapter
- 10 or 40GbE networks

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

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



Viewing and Configuring Cisco usNIC using the Cisco IMC GUI

Before You Begin

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

Procedure

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

Name	Description
Name	The name for the vNIC that is the parent of the usNIC. Note This field is read-only.
usNIC field	The number of usNICs assigned to the specific vNIC. Enter an integer between 0 and 225. To assign additional usNICs to a specified vNIC, enter value higher than the existing value. To delete usNICs from a specified vNIC, enter value smaller than the existing value. To delete all the usNICs assigned to a vNIC, enter zero.
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.
Completion Queue Count field	The number of completion queue resources to allocate. In general, the number of completion queue resources you should allocate is equal to the number of transmit queue resources plus the number of receive queue resources. Enter an integer between 1 and 512.
Transmit Queue Ring Size field	The number of descriptors in each transmit queue. Enter an integer between 64 and 4096.
Receive Queue Ring Size field	The number of descriptors in each receive queue. Enter an integer between 64 and 4096.

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

Name	Description
TCP Rx Checksum check box	If checked, the CPU sends all packet checksums to the hardware for validation. This option may reduce CPU overhead. If cleared, the CPU validates all packet checksums.

- Step 7** Click **Save Changes**.
The changes take effect upon the next server reboot.

Viewing usNIC Properties

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Host Ethernet Interfaces** pane's **usNIC Properties** area, review the information in the following fields:

Name	Description
Name	The name for the vNIC that is the parent of the usNIC. Note This field is read-only.
usNIC field	The number of usNICs assigned to the specific vNIC. Enter an integer between 0 and 225. To assign additional usNICs to a specified vNIC, enter value higher than the existing value. To delete usNICs from a specified vNIC, enter value smaller than the existing value. To delete all the usNICs assigned to a vNIC, enter zero.
Transmit Queue Count field	The number of transmit queue resources to allocate. Enter an integer between 1 and 256.
Receive Queue Count field	The number of receive queue resources to allocate. Enter an integer between 1 and 256.

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>
Transmit Queue Ring Size field	<p>The number of descriptors in each transmit queue.</p> <p>Enter an integer between 64 and 4096.</p>
Receive Queue Ring Size field	<p>The number of descriptors in each receive queue.</p> <p>Enter an integer between 64 and 4096.</p>
Interrupt Count field	<p>The number of interrupt resources to allocate. In general, this value should be equal to the number of completion queue resources.</p> <p>Enter an integer between 1 and 514.</p>
Interrupt Coalescing Type drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • MIN—The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • IDLE—The system does not send an interrupt until there is a period of no activity lasting as least as long as the time specified in the Coalescing Time field.
Interrupt Coalescing Timer Time field	<p>The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent.</p> <p>Enter an integer between 1 and 65535. To turn off interrupt coalescing, enter 0 (zero) in this field.</p>
Class of Service field	<p>The class of service to associate with traffic from this usNIC.</p> <p>Select an integer between 0 and 6, with 0 being lowest priority and 6 being the highest priority.</p> <p>Note This option cannot be used in VNTAG mode.</p>

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

Configuring iSCSI Boot Capability

Configuring iSCSI Boot Capability for vNICs

When the rack-servers are configured in a standalone mode, and when the VIC adapters are directly attached to the Nexus 5000 and Nexus 6000 family of switches, you can configure these VIC adapters to boot the servers remotely from iSCSI storage targets. You can configure Ethernet vNICs to enable a rack server to load the host OS image from remote iSCSI target devices.

To configure the iSCSI boot capability on a vNIC:

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



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

Configuring iSCSI Boot Capability on a vNIC

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

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Ethernet Interfaces** area, select the **iSCSI Boot Properties** area.
- Step 5** In the **General Area**, update the following fields:

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

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

Step 6 In the **Initiator Area**, update the following fields:

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

Step 7 In the **Primary Target Area**, update the following fields:

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

Step 8 In the **Secondary Target Area**, update the following fields:

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

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

Step 9 Click **Save Changes**.

Removing iSCSI Boot Configuration from a vNIC

Before You Begin

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

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** In the **Adapter Card** pane, click the **vNICs** tab.
- Step 3** In the vNICs pane, click **eth0** or **eth1**.
- Step 4** In the **Ethernet Interfaces** area, select the **iSCSI Boot Properties** area.
- Step 5** Click the **Unconfigure iSCSI** button at the bottom of the area.
-

Backing Up and Restoring the Adapter Configuration

Exporting the Adapter Configuration

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

- TFTP
- FTP
- SFTP
- SCP
- HTTP

Before You Begin

Obtain the remote server IP address.

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** Click the **Adapter Card** tab.
The **General** tab appears.
- Step 3** In the **Actions** area of the **General** tab, click **Export Configuration**.

The **Export Adapter Configuration** dialog box opens.

Step 4 In the **Export Adapter Configuration** dialog box, update the following fields:

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

Step 5 Click **Export Configuration**.

Importing the Adapter Configuration

Procedure

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

Step 2 Click the **Adapter Card** tab.
The **General** tab appears.

Step 3 In the **Actions** area of the **General** tab, click **Import Configuration**.
The **Import Adapter Configuration** dialog box opens.

Step 4 In the **Import Adapter Configuration** dialog box, update the following fields:

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

Step 5 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 **Networking** menu.
- Step 2** Click the **Adapter Card** tab.
The **General** tab appears.
- Step 3** In the **Actions** area of the **General** tab, click **Reset To Defaults** and click **OK** to confirm.
- Note** Resetting the adapter to default settings sets the port speed to 4 X 10 Gbps. Choose 40 Gbps as the port speed only if you are using a 40 Gbps switch.
-

Resetting the Adapter

Procedure

- Step 1** In the **Navigation** pane, click the **Networking** menu.
- Step 2** Click the **Adapter Card** tab.
The **General** tab appears.
- Step 3** In the **Actions** area of the **General** tab, click **Reset** and click **Yes** to confirm.
- Note** Resetting the adapter also resets the host and requires a reformat.
-

