

Cisco usNIC Deployment Guide for Cisco UCS C-Series Rack-Mount Standalone Servers

First Published: 2016-01-20

Last Modified: 2016-09-17

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 adapter, 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) or Libfabric interface instead of sockets or other communication APIs.

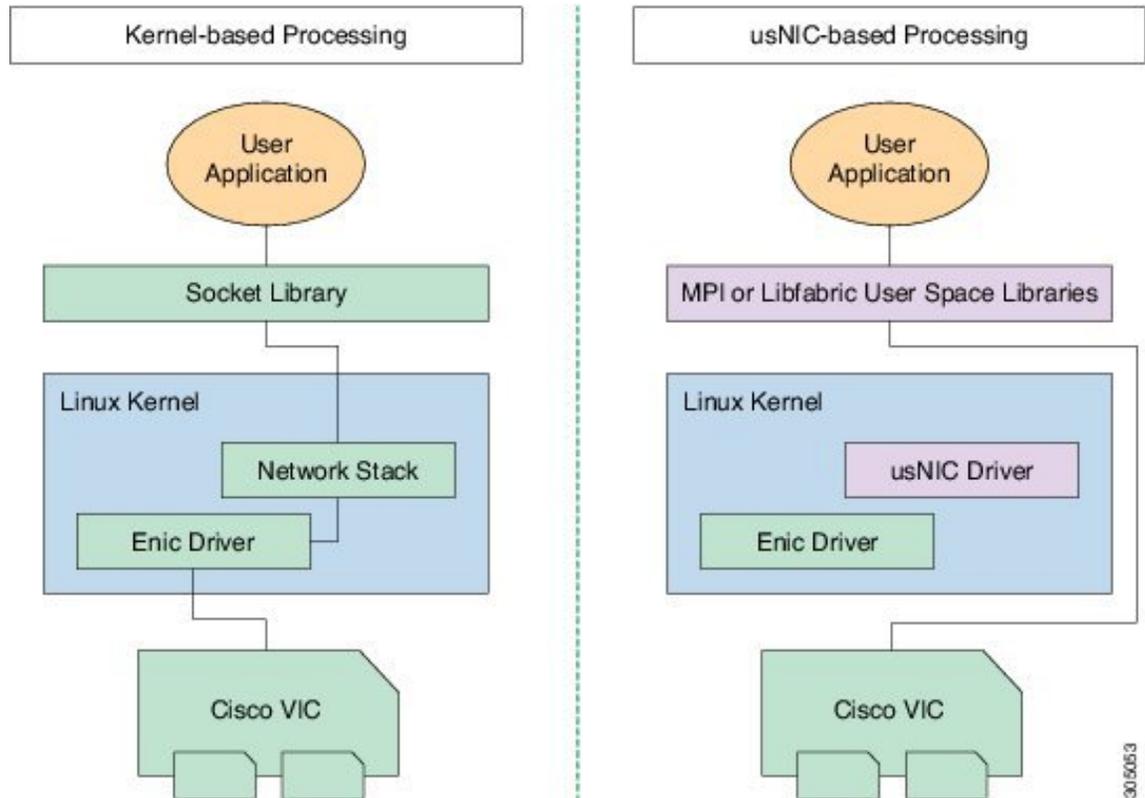
Cisco usNIC offers the following benefits for your applications:

- Provides a low-latency and high-throughput communication transport.
- Employs the standard and application-independent Ethernet protocol.
- Low jitter, or near constant latency, communications.
- Takes advantage of low latency 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 usNIC.

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



Cisco usNIC Prerequisites

To benefit from Cisco usNIC, your configuration has the following prerequisites:

- UCS Driver ISO. For more information, see [Downloading Cisco UCS VIC drivers](#).
- A supported Linux operating system distribution release. For more information, see the appropriate [Hardware and Software Interoperability guide](#).
- A supported MPI implementation, such as the Cisco Open MPI distribution (included on the Cisco UCS Driver ISO), or version 4 or 5 of the Intel[®] MPI Library. If the Intel[®] MPI Library is used, the network must be configured with flow control enabled.

Configuring Cisco usNIC



Note

The Cisco usNIC packages do not support the upgrade or downgrade of an operating system. To update the operating system, first ensure you uninstall the usNIC packages, update your operating system, and then reinstall the usNIC drivers.

Alternatively, you can update the operating system, uninstall the usNIC drivers, and then reinstall the usNIC drivers.

Before You Begin

Make sure that the following software and hardware components are installed on the Cisco UCS server:

- A supported Linux operating system distribution release. For more information, see the appropriate [Hardware and Software Interoperability guide](#).
- GCC, G++, and Gfortran
- DAT user library (if using Intel[®] MPI)
- libnl user library (either version 1 or version 3)
- Cisco UCS VIC second generation or later adapter
- rdma package



Important

For information on supported Linux operating system distributions, see the content of the usNIC folder that is included in the UCS Drivers ISO bundle. See [Cisco UCS Virtual Interface Card Drivers for Linux Installation Guide](#).

Step 1

Ensure the kernel option **CONFIG_INTEL_IOMMU** is selected in the kernel. Enable the Intel IOMMU driver in the Linux kernel by manually adding 'intel_iommu=on' in the grub.conf file (/boot/grub/grub.conf) for GRUB based systems: For example, if your grub.conf file contains a "kernel" line such as `kernel (hd0,0)/vmlinuz LANG=en_US.UTF-8 KEYTABLE=us`, then you will add 'intel_iommu=on' to the end as shown below:

```
kernel (hd0,0)/vmlinuz LANG=en_US.UTF-8 KEYTABLE=us intel_iommu=on
```

For grub2 system, you can use the following:

```
edit /etc/grub2.cfg, add setting to kernel args, run grub2-mkconfig -o/boot/grub2/grub.cfg
```

Step 2

Reboot your Cisco UCS server.

You must reboot your server for the changes to take effect after you configure Cisco usNIC.

Step 3

Verify that the running kernel has booted with the intel_iommu=on option.

```
$ cat /proc/cmdline | grep iommu
```

Step 4

Install the Cisco usNIC Linux drivers.

For more information about installing the drivers, see the section on "Installing Linux Software Packages for Cisco usNIC" in this guide.

What to Do Next

After you complete configuring Cisco usNIC and installing the Linux drivers, verify that Cisco usNIC is functioning properly. For more information about how to verify the installation, see the section on "Verifying the Cisco usNIC Installation for Cisco UCS C-Series Rack-Mount Standalone Servers" in this guide.

Configuring Cisco usNIC Using the CIMC GUI



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

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

Before You Begin

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

-
- Step 1** Log into the CIMC GUI.
For more information about how to log into CIMC, see [Cisco UCS C-Series Servers Integrated Management Controller GUI Configuration Guide](#).
- Step 2** In the **Navigation** pane, click the **Server** tab.
- Step 3** On the **Server** tab, click **Inventory**.
- Step 4** In the **Inventory** pane, click the **Cisco VIC Adapters** tab.
- Step 5** In the **Adapter Cards** area, select the adapter card.
If the server is powered on, the resources of the selected adapter card appear in the tabbed menu below the **Adapter Cards** area.
- Step 6** In the tabbed menu below the **Adapter Cards** area, click the **vNICs** tab.
- Step 7** In the **Host Ethernet Interfaces** area, select a vNIC from the table.
Note For each vNIC that you want to configure as a usNIC, select the vNIC entry from the table and specify its properties as explained in steps 9 through step 18.
- Step 8** Click **usNIC** to open the **usNIC Properties** dialog box.
- Step 9** In the **usNICs** property, specify the number of Cisco usNICs that you want to create.
Each MPI process that is running on the server requires a dedicated usNIC. You might need to create up to 64 usNICs to sustain 64 MPI processes running simultaneously. We recommend that you create at least as many usNICs, per usNIC-enabled vNIC, as the number of physical cores on your server. For example, if you have 8 physical cores on your server, create 8 usNICs.

Step 10 In the **Properties** area, update the following fields:

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

Step 11 Click **Apply**.

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

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

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

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

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

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

Step 17 Click **Save Changes**.

The changes take effect upon the next server reboot.

Creating Cisco usNIC Using the CIMC CLI



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

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

Before You Begin

You must log in to the CIMC CLI with administrator privileges to perform this task.

DETAILED STEPS

	Command or Action	Purpose
Step 1	server# scope chassis	Enters chassis command mode.
Step 2	server/chassis# scope adapter <i>index</i>	Enters the command mode for the adapter card at the PCI slot number specified by <i>index</i> . Note Make sure that the server is powered on before you attempt to view or change adapter settings. To view the index of the adapters configured on your server, use the show adapter command.
Step 3	server/chassis/adapter# scope host-eth-if { eth0 eth1 }	Enters the command mode for the vNIC. Specify the Ethernet ID based on the number of vNICs that you have configured in your environment. For example, specify eth0 if you configured only one vNIC.
Step 4	server/chassis/adapter/host-eth-if# create usnic-config 0	Creates a usNIC config and enters its command mode. Make sure that you always set the index value to 0. Note To create a Cisco usNIC for the first time for a given vNIC using the CLI, you must first create a usnic-config . Subsequently, you only need to scope into the usnic-config and modify the properties for Cisco usNIC. For more information about modifying Cisco usNIC properties, see <i>Modifying a Cisco usNIC value using the CLI</i> .
Step 5	server/chassis/adapter/host-eth-if/usnic-config# set cq-count <i>count</i>	Specifies the number of completion queue resources to allocate. We recommend that you set this value to 6. The number of completion queues equals the number of transmit queues plus the number of receive queues.
Step 6	server/chassis/adapter/host-eth-if/usnic-config# set rq-count <i>count</i>	Specifies the number of receive queue resources to allocate. We recommend that you set this value to 6.
Step 7	server/chassis/adapter/host-eth-if/usnic-config# set tq-count <i>count</i>	Specifies the number of transmit queue resources to allocate. We recommend that you set this value to 6.
Step 8	server/chassis/adapter/host-eth-if/usnic-config# set usnic-count <i>number of usNICs</i> .	Specifies the number of Cisco usNICs to create. Each MPI process that is running on the server requires a dedicated Cisco usNIC. Therefore, you might need to create up to 64 Cisco usNICs to sustain 64 MPI processes running simultaneously. We recommend that you create at least as many Cisco usNICs, per Cisco usNIC-enabled vNIC, as the number of physical cores on your server. For example, if you have 8 physical cores on your server, create 8 Cisco usNICs.
Step 9	server/chassis/adapter/host-eth-if/usnic-config# commit	Commits the transaction to the system configuration. Note The changes take effect when the server is rebooted.

	Command or Action	Purpose
Step 10	server/chassis/adapter/host-eth-if/usnic-config# exit	Exits to host Ethernet interface command mode.
Step 11	server/chassis/adapter/host-eth-if# exit	Exits to adapter interface command mode.
Step 12	server/chassis/adapter# exit	Exits to chassis interface command mode.
Step 13	server/chassis# exit	Exits to server interface command mode.
Step 14	server# scope bios	Enters Bios command mode.
Step 15	server/bios# scope advanced	Enters the advanced settings of BIOS command mode.
Step 16	server/bios/advanced# set IntelVTD Enabled	Enables the Intel Virtualization Technology.
Step 17	server/bios/advanced# set ATS Enabled	Enables the Intel VT-d Address Translation Services (ATS) support for the processor.
Step 18	server/bios/advanced# set CoherencySupport Enabled	Enables Intel VT-d coherency support for the processor.
Step 19	server /bios/advanced# commit	Commits the transaction to the system configuration. Note The changes take effect when the server is rebooted.

This example shows how to configure Cisco usNIC properties:

```

Server # scope chassis
server /chassis # show adapter
server /chassis # scope adapter 2
server /chassis/adapter # scope host-eth-if eth0
server /chassis/adapter/host-eth-if # create usnic-config 0
server /chassis/adapter/host-eth-if/usnic-config *# set usnic-count 64
server /chassis/adapter/host-eth-if/usnic-config *# set cq-count 6
server /chassis/adapter/host-eth-if/usnic-config *# set rq-count 6
server /chassis/adapter/host-eth-if/usnic-config *# set tq-count 6
server /chassis/adapter/host-eth-if/usnic-config *# commit
Committed settings will take effect upon the next server reset
server /chassis/adapter/host-eth-if/usnic-config # exit
server /chassis/adapter/host-eth-if # exit
server /chassis/adapter # exit
server /chassis # exit
server # exit
server# scope bios
server /bios # scope advanced
server /bios/advanced # set IntelVTD Enabled
server /bios/advanced *# set ATS Enabled*
server /bios/advanced *# set CoherencySupport Enabled
server /bios/advanced *# commit
Changes to BIOS set-up parameters will require a reboot.
Do you want to reboot the system?[y|N]y
A system reboot has been initiated.

```

Modifying a Cisco usNIC value using the CIMC CLI

Before You Begin

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>server# scope chassis</code>	Enters chassis command mode.
Step 2	<code>server/chassis# scope adapter <i>index</i></code>	Enters the command mode for the adapter card at the PCI slot number specified by <i>index</i> . Note Make sure that the server is powered on before you attempt to view or change adapter settings. To view the index of the adapters configured on your server, use the show adapter command.
Step 3	<code>server/chassis/adapter# scope host-eth-if {eth0 eth1}</code>	Enters the command mode for the vNIC. Specify the Ethernet ID based on the number of vNICs that you have configured in your environment. For example, specify eth0 if you configured only one vNIC.
Step 4	<code>server/chassis/adapter/host-eth-if# scope usnic-config 0</code>	Enters the command mode for the usNIC. Make sure that you always set the index value as 0 to configure a Cisco usNIC.
Step 5	<code>server/chassis/adapter/host-eth-if/usnic-config# set usnic-count <i>number of usNICs</i> .</code>	Specifies the number of Cisco usNICs to create. Each MPI process running on the server requires a dedicated Cisco usNIC. Therefore, you might need to create up to 64 Cisco usNIC to sustain 64 MPI processes running simultaneously. We recommend that you create at least as many Cisco usNIC, per Cisco usNIC-enabled vNIC, as the number of physical cores on your server. For example, if you have 8 physical cores on your server, create 8 usNICs.
Step 6	<code>server /chassis/adapter/host-eth-if/usnic-config# commit</code>	Commits the transaction to the system configuration. Note The changes take effect when the server is rebooted.
Step 7	<code>server/chassis/adapter/host-eth-if/usnic-config# exit</code>	Exits to host Ethernet interface command mode.
Step 8	<code>server/chassis/adapter/host-eth-if# exit</code>	Exits to adapter interface command mode.
Step 9	<code>server/chassis/adapter# exit</code>	Exits to chassis interface command mode.
Step 10	<code>server/chassis# exit</code>	Exits to server interface command mode.

This example shows how to configure Cisco usNIC properties:

```
server # scope chassis
server /chassis # show adapter
```

```

server /chassis # scope adapter 2
server /chassis/adapter # scope host-eth-if eth0
server /chassis/adapter/host-eth-if # scope usnic-config 0
server /chassis/adapter/host-eth-if/usnic-config # set usnic-count 32
server /chassis/adapter/host-eth-if/usnic-config # commit
Committed settings will take effect upon the next server reset
server /chassis/adapter/host-eth-if/usnic-config # exit
server /chassis/adapter/host-eth-if # exit
server /chassis/adapter # exit
server /chassis # exit
server # exit

```

Deleting Cisco usNIC from a vNIC

Before You Begin

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	server# scope chassis	Enters chassis command mode.
Step 2	server/chassis# scope adapter index	Enters the command mode for the adapter card at the PCI slot number specified by <i>index</i> . Note Make sure that the server is powered on before you attempt to view or change adapter settings. To view the index of the adapters configured on you server, use the show adapter command.
Step 3	server/chassis/adapter# scope host-eth-if {eth0 eth1}	Enters the command mode for the vNIC. Specify the Ethernet ID based on the number of vNICs that you have configured in your environment. For example, specify eth0 if you configured only one vNIC.
Step 4	Server/chassis/adapter/host-eth-if# delete usnic-config 0	Deletes the Cisco usNIC configuration for the vNIC.
Step 5	Server/chassis/adapter/host-eth-if# commit	Commits the transaction to the system configuration Note The changes take effect when the server is rebooted.

This example shows how to delete the Cisco usNIC configuration for a vNIC:

```

server # scope chassis
server/chassis # show adapter
server/chassis # scope adapter 1
server/chassis/adapter # scope host-eth-if eth0
server/chassis/adapter/host-eth-if # delete usnic-config 0
server/chassis/host-eth-if/iscsi-boot *# commit
New host-eth-if settings will take effect upon the next adapter reboot

server/chassis/host-eth-if/usnic-config #

```

Installing Linux Software Packages for Cisco usNIC

The following section lists the content of the usNIC folder, specific for each supported Linux operating system distribution that is included in the UCS Drivers ISO bundle. Documentation about known issues and installation instructions are also included in the README file in the usNIC folder.

- **kmod-usnic_verbs-{version}.x86_64.rpm**—Linux kernel verbs driver for the usNIC feature of the Cisco VIC SR-IOV Ethernet NIC.
- **libdaplusnic_verbs-{version}.x86_64.rpm**— User space library DAPL plugin for usNIC.
- **openmpi-cisco-{version}.x86_64.rpm**—Cisco usNIC Open MPI — Open MPI with the Cisco usNIC BTL MPI transport.
- **usnic_tools-{version}.x86_64.rpm** — Utility programs for usNIC.
- **libfabric-cisco-{version}.x86_64.rpm**— Libfabric package with built-in Cisco usNIC transport.
- **libusnic_verbs-{version}.x86_64.rpm**— A dummy library that causes the libibverbs library to skip Cisco usNIC Linux devices (because Cisco usNIC functionality is exposed through libfabric, not libibverbs).
- **usd_tools-{version}.x86_64.rpm**—Scripts that help verifying system environments running Cisco usNIC applications and collecting tech support information.

Before You Begin

Make sure that you have configured the Cisco usNIC properties in CIMC. For more information about how to configure the properties, see *Configuring Cisco usNIC*.

You must also make sure that the host OS distribution on which you want to install Cisco usNIC has a supported version of the Cisco enic driver installed. The Cisco enic driver is the Linux kernel networking driver for the Cisco VIC SR-IOV Ethernet NIC.

-
- Step 1** Upgrade to the latest version of the enic driver included in the Cisco UCS ISO for your Linux distribution.
- Step 2** Install the Cisco usNIC software packages from the Cisco UCS Drivers ISO for your Linux distribution.
- Step 3** **# chkconfig rdma on**
Enables Linux RDMA services. Once enabled, RDMA services will be started automatically after a system reboot.
- Note** You may need to perform this step on some Linux operating systems distributions, such as RHEL 6.4.
- Step 4** Reboot your server for the installation changes to take effect automatically.
- Important** If you do not want to reboot your server, you can manually load the kernel modules to ensure the system loads the correct version of the driver and enforces the new memory lock configurations. For more information about how to load the modules, see *Manually Loading the Kernel Modules for Cisco usNIC*.
-

Source code for Linux Cisco usNIC software packages

The source code for the Cisco usNIC software packages is provided on the Cisco UCS Drivers ISO. It is recommended that you do not mix the source code and binary package installations.

Manually Loading the Kernel Modules for Cisco usNIC

If you do not want to reboot your server, you can manually load the Cisco usNIC kernel modules by using the following steps.

Before You Begin

Ensure you delete all the existing versions of the driver before you load the latest version of the driver. This will help you configure the system successfully.

DETAILED STEPS

	Command or Action	Purpose
Step 1	# <code>rmmod enic</code>	Unloads the existing enic driver module. Note Make sure that you are not logged into the OS using the network, for example, via SSH. Otherwise, your network connection might get permanently disconnected. Alternatively, you can log in to the server using the KVM to perform this step.
Step 2	# <code>modprobe enic</code>	Loads the enic driver module.
Step 3	# <code>modprobe usnic_verbs</code>	Loads the usnic_verbs driver module.

Upgrading the Linux Software Packages for Cisco usNIC

-
- Step 1** Follow the procedure in "Uninstalling Linux Software Packages for Cisco usNIC" to uninstall the previous versions of the usNIC software packages.
- Step 2** Follow the procedure in "Installing Linux Software Packages for Cisco usNIC" to install usNIC software packages from the Cisco UCS Drivers ISO for your Linux distribution.
-

Uninstalling Linux Software Packages for Cisco usNIC

Step 1 Uninstall the following usNIC software packages:

- usnic_tools
- openmpi-cisco
- libdaplusnic
- libusnic_verbs
- kmod-usnic_verbs

- libfabric-cisco-devel
- libfabric-cisco
- usd_tools

Step 2 Reboot your Cisco server.

Adding MPI to User Environments

Before MPI applications can be compiled and launched, an MPI implementation must be added to each user's environment. It is recommended that you only add one MPI implementation to a user's environment at a time.

Environment for the Cisco Open MPI

For Cisco Open MPI, two scripts are installed by the openmpi-cisco software package to help set the required environment variables. One script is for Bourne shell users; the other script is for C shell users:

- `/opt/cisco/openmpi-vars.sh` (For Bourne shell and Bourne again shell users)
- `/opt/cisco/openmpi-vars.csh` (For C shell users)

The appropriate script should be sourced as part of the users's shell startup / login sequence (even for non-interactive shells).

Environment for the Intel[®] MPI Library

In addition to the instructions provided by the Intel[®] MPI Library documentation, additional environment variables must be set in each user's environment to enable Cisco usNIC functionality. Two scripts are installed by the libdaplusnic software package to help set the required environment variables. One script is for Bourne shell users; the other script is for C shell users:

- `/opt/cisco/intelmpi-usnic-vars.sh` (For Bourne shell and Bourne again shell users)
- `/opt/cisco/intelmpi-usnic-vars.csh` (For C shell users)

The appropriate script should be sourced as part of the users's shell startup / login sequence.

Using the Intel[®] MPI Library with usNIC requires the network to be configured with flow control enabled. This can be either IEEE 802.3x link-level flow control or IEEE 802.1Qbb Priority-based Flow Control (PFC). This feature is sometimes also called "no-drop." Refer to the configuration guide for the switch(es) in your network for information about enabling flow control. If flow control is not enabled in the network, then applications utilizing the Intel[®] MPI Library may work correctly, but possibly with extremely degraded performance.

In deployments of the Intel[®] MPI Library, the MPI traffic must have flow control enabled on all Cisco usNIC ports, and no-drop or platinum QoS system Class with default CoS value 5 in Cisco UCS Manager. Please refer to the *Cisco UCS Manager Network Management guide* section on "Quality of Service".

Adding Libfabric to User Environments

For Libfabric, two scripts are installed by the libfabric package to help set the required environment variables. One script is for the Bourne shell users; the other script is for the C shell users:

- /opt/cisco/libfabric-vars.sh
- /opt/cisco/libfabric-vars.csh

The appropriate script should be sourced as part of the users's shell startup / login sequence.

Verifying the Cisco usNIC Installation for Cisco UCS C-Series Rack-Mount Standalone Servers

After you install the required Linux drivers for Cisco usNIC, perform the following procedure at the Linux prompt to make sure that the installation completed successfully.



Note

The examples shown below are configurations verified on Linux operating system distribution RHEL 6.5.

Step 1

Search and verify if the **usnic_verbs** kernel module was loaded during the OS driver installation.

```
$ lsmod | grep usnic_verbs
```

The following details are displayed when you enter the **lsmod | grep usnic_verbs** command. The kernel modules listed on your console may differ based on the modules that you have currently loaded in your OS.

Important usnic_verbs is displayed in the output, showing that this kernel module is loaded and active in the kernel.

```
usnic_verbs          73762  2
ib_core              74355  11
ib_ipoib,rdma_ucm,ib_ucm,ib_uverbs,ib_umad,rdma_cm,ib_cm,iw_cm,ib_sa,ib_mad,usnic_verbs
enic                 73723  1 usnic_verbs
```

Step 2

View the configuration of Cisco usNIC-enabled NICs.

```
$ /opt/cisco/usnic/bin/usnic_devinfo
```

The following section is a brief example of the results that are displayed when you execute the **usnic_devinfo** command. The results may differ based on your current installation. When the results are displayed on your console, make sure that the correct corresponding Ethernet interface is shown for each usNIC interface, that the link state is UP, and that the number of VFs, CQs per VF, and QPs per VF are all the values that you set in the "Create Cisco usNIC using the CIMC GUI" or "Creating Cisco usNIC using the CIMC CLI" sections.

The following example shows two interfaces. (**usnic_1** and **usnic_0**) that are configured on a Cisco UCS VIC adapter. If you configured only one Cisco usNIC-enabled vNIC, you will see a listing for only **usnic_0**.

```
Rack:
usnic_0:
Interface:          eth0
MAC Address:       58:ac:78:f2:84:e8
IP Address:        10.10.10.2
Netmask:          255.255.255.0
Prefix len:       24
MTU:              9000
Link State:       UP
Bandwidth:        40 Gb/s
Device ID:        UCSC-PCIE-C40Q-03 [VIC 1385] [0x014d]
Vendor ID:        4407
Vendor Part ID:   207
```

```

Firmware:          4.1(1d)
VFs:              64
CQ per VF:        6
QP per VF:        6
Interrupts per VF: 6
Max CQ:           384
Max CQ Entries:   65535
Max QP:           384
Max Send Credits: 4095
Max Recv Credits: 4095
Capabilities:
  Map per res:    yes
  PIO sends:      yes
  CQ interrupts:  no
usnic_1:
Interface:        eth1
MAC Address:      58:ac:78:f2:84:e9
IP Address:       10.20.20.2
Netmask:         255.255.255.0
Prefix len:      24
MTU:             9000
Link State:      UP
Bandwidth:       40 Gb/s
Device ID:       UCSB-PCIE-C40Q-03 [VIC 1385] [0x014d]
Vendor ID:       4407
Vendor Part ID:  207
Firmware:        4.1(1d)
VFs:             64
CQ per VF:       6
QP per VF:       6
Interrupts per VF: 6
Max CQ:          384
Max CQ Entries:  65535
Max QP:          384
Max Send Credits: 4095
Max Recv Credits: 4095
Capabilities:
  Map per res:    yes
  PIO sends:      yes
  CQ interrupts:  no

Blade:
usnic_0:
Interface:        eth0
MAC Address:      00:25:b5:01:01:6e
IP Address:       10.10.10.2
Netmask:         255.255.255.0
Prefix len:      24
MTU:             9000
Link State:      UP
Bandwidth:       10 Gb/s
Device ID:       UCSB-VIC-M83-8P [VIC 1380] [0x0137]
Vendor ID:       4407
Vendor Part ID:  207
Firmware:        4.1(1d)
VFs:             32
CQ per VF:       6
QP per VF:       6
Interrupts per VF: 6
Max CQ:          192
Max CQ Entries:  65535
Max QP:          192
Max Send Credits: 4095
Max Recv Credits: 4095
Capabilities:
  Map per res:    yes
  PIO sends:      yes
  CQ interrupts:  no
usnic_1:
Interface:        eth1
MAC Address:      00:25:b5:01:01:9e

```

```

IP Address:          10.20.20.2
Netmask:            255.255.255.0
Prefix len:        24
MTU:               9000
Link State:        UP
Bandwidth:         10 Gb/s
Device ID:         UCSB-VIC-M83-8P [VIC 1380] [0x0137]
Vendor ID:         4407
Vendor Part ID:    207
Firmware:          4.1(1d)
VFs:              32
CQ per VF:        6
QP per VF:        6
Interrupts per VF: 6
Max CQ:           192
Max CQ Entries:   65535
Max QP:           192
Max Send Credits: 4095
Max Recv Credits: 4095
Capabilities:
  Map per res:     yes
  PIO sends:       yes
  CQ interrupts:   no

```

Step 3 Run the `usnic_check` script to view the installed RPMs and their versions.

```
$ /opt/cisco/usnic/bin/usnic_check
```

If any errors occurred during the OS driver installation, warnings are generated.

If the `usnic_verbs` module failed to load, the following brief example shows the warnings that are generated:

```

$ /opt/cisco/usnic/bin/usnic_check
enic RPM version 2.3.0.18 installed
usnic_verbs RPM version 1.0.4.416.rhel6u5 installed
WARNING: usnic_verbs module not loaded
libdaplusnic RPM version 2.0.39cisco1.0.0.401 installed
libfabric RPM version 1.1.1cisco1.0.0.410.rhel6u5 installed
libusnic_verbs RPM version 2.0.0.406.rhel6u5 installed
Using /opt/cisco/openmpi/bin/ompi_info to check Open MPI info...
Open MPI version 1.10.1cisco1.0.0.410.rhel6u5 installed
WARNING: No usnic devices found
WARNING: No usnic devices found
3 warnings

```

Step 4 Verify that the Cisco usNIC network packets are being transmitted correctly between the client and server hosts.

a) Determine the name of the Ethernet interface associated with the Cisco usNIC on the server host.

```
[server]$ /opt/cisco/usnic/bin/usnic_status
usnic_0: 0000:06:0.0, eth11, 04:62:73:c5:3a:a6, 64 VFs
Per VF: 6 WQ, 6 RQ, 6 CQ, 4 INT
```

```

In use:
0 VFs, 0 QPs, 0 CQs

```

```

usnic_1: 0000:07:0.0, eth10, 04:62:73:c5:3a:a7, 64 VFs
Per VF: 6 WQ, 6 RQ, 6 CQ, 4 INT

```

```

In use:
0 VFs, 0 QPs, 0 CQs

```

b) Determine the IP address for the Ethernet interface.

```
<server>$ ip addr show dev eth4 | grep "inet[^\"]"
inet 10.1.0.1/16 brd 10.1.255.255 scope global eth4
```

c) Run the `usd_pingpong` program on the server host.

```
[server]$ /opt/cisco/usnic/bin/usd_pingpong
```

For more information about the command line options used with the **usd_pingpong** program, see the output of `usd_pingpong --help`.

- d) Execute the **usd_pingpong** program on the client host by using the IP address that corresponds to the Cisco usNIC on the server host.

```
[client]# /opt/cisco/usnic/bin/usd_pingpong -h SERVER_IP_ADDRESS
```

The following example shows the results that are displayed when you run the **usd_pingpong** program.

```
Server-side:
/opt/cisco/usnic/bin/usd_pingpong -d usnic_0
open usnic_0 OK, IP=10.10.10.14
QP create OK, addr -h 10.10.10.14 -p 3333
Waiting for setup...

Client-side
/opt/cisco/usnic/bin/usd_pingpong -d usnic_0 -h 10.10.10.14
open usnic_0 OK, IP=10.10.10.24
QP create OK, addr -h 10.10.10.24 -p 3333
sending params...
payload_size=4, pkt_size=46
posted 63 RX buffers, size=64 (4)
100000 pkts, 4.137 us / HRT
```

- Step 5** Download, compile, and execute the **ring_c** test program to validate that the MPI traffic is correctly transmitted between the client and server hosts.

You can obtain the **ring_c** test program from this link: https://raw.githubusercontent.com/open-mpi/ompi-release/v1.8/examples/ring_c.c.

The following example shows how to use the **wget** utility to obtain, compile, and execute the **ring_c**. Alternatively, you can use other methods of obtaining and running the test program.

Note Run the following commands with a single MPI implementation setup in your environment.

```
$ wget --no-check-certificate
https://raw.githubusercontent.com/open-mpi/ompi-release/v1.8/examples/ring_c.c
--2015-04-23 10:11:42-- https://raw.githubusercontent.com/open-mpi/ompi-release/v1.8/examples/ring_c.c
Resolving raw.githubusercontent.com... 199.27.74.133
Connecting to raw.githubusercontent.com|199.27.74.133|:443... connected.
WARNING: certificate common name \u201cwww.github.com\u201d doesn\u2019t match requested host name
\u201craw.githubusercontent.com\u201d.
HTTP request sent, awaiting response... 200 OK
Length: 2418 (2.4K) [text/plain]

$ mpicc ring_c.c -o ring_c
[no output]

$ mpiexec --host host1,host2 -n 4 ./ring_c
Process 0 sending 10 to 1, tag 201 (4 processes in ring)
Process 0 sent to 1
Process 0 decremented value: 9
Process 0 decremented value: 8
Process 0 decremented value: 7
Process 0 decremented value: 6
Process 0 decremented value: 5
Process 0 decremented value: 4
Process 0 decremented value: 3
Process 0 decremented value: 2
Process 0 decremented value: 1
Process 0 decremented value: 0
Process 0 exiting
Process 2 exiting
Process 1 exiting
Process 3 exiting ...
```

Note If desired, setup a different MPI implementation in your environment and re-run the **mpicc** and **mpirun** commands to verify that MPI implementation with Cisco usNIC functionality.

If the **usd_pingpong** program and the **ring_c** program executed successfully, you should now be able to run MPI applications over Cisco usNIC.

Troubleshooting Information

Problem

Viewing the list of installed RPMs using **usnic_check** causes the following:

- 1 A warning such as **No usnic devices found**.
- 2 A version mismatch error such as **usnic_verbs_xxxx does not match installed version**.

Possible Cause

A previously installed version can cause this error.

Solution

- 1 List all the installed versions using the following command: **rpm -qa|grep usnic_verbs**
- 2 Uninstall all versions using the following command: **rpm -e**
- 3 Make sure that the module has been removed.
- 4 Re-install all the RPMs.

Problem

Verifying that Cisco usNIC packets are being transmitted correctly between client and server using **usd_pingpong** causes the following errors:

- 1 “No such address or device” error. See the following example:

```
# /opt/cisco/usnic/bin/usd_pingpong
usd_open: No such device or address
```

- 2 The server gets stuck at “Waiting for setup...” stage. See the following example:

```
# /opt/cisco/usnic/bin/usd_pingpong -d usnic_0
open usnic_0 OK, IP=10.43.10.1
QP create OK, addr -h 10.43.10.1 -p 3333
Waiting for setup...
```

Possible Cause

- 1 The Cisco usNIC connection policy is not assigned or set as 'not set' in the vNIC interface.
- 2 The server side does not receive packets from the client side.

Solution

- 1 Make sure that valid Cisco usNIC connection policy is configured in usNIC Connection Policies and assigned to the vNICs in the service profile.
- 2 Make sure that IP addresses of the Cisco usNIC devices on both the server and client are configured correctly.
- 3 Make sure that the client pingpong is attempting to send packets to the correct server IP address of Cisco usNIC device.

Problem

Running the Cisco usNIC traffic using the `mpirun` causes the following errors:

MTU size mismatch error. See the following example:

Example:

```
$ mpirun --host node05,node06 -np 12 --mca btl usnic,sm,self --mca btl_usnic_if_include usnic_1 IMB-MPI1 Sendrecv
```

The MTU does not match on local and remote hosts. All interfaces on all hosts participating in an MPI job must be configured with the same MTU. The usNIC interface listed below will not be used to communicate with this remote host.

```
Local host:      node05
usNIC interface: usnic_1
Local MTU:      8958
Remote host:    node06
Remote MTU:    1458
```

Possible Cause

- 1 The MTU size is incorrectly set on the appropriate VLANs.
- 2 The MTU size is incorrectly set in the QoS.

Solution

Make sure that the MTU size has been set correctly on the VLANs and QoS.

See: [Configuring QoS System Classes with the LAN Uplinks Manager](#).

Problem

Installing a Cisco enic driver causes the following Cisco enic dependency errors:

```
# rpm -ivh kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64.rpm
error: Failed dependencies:
        ksym(enic_api_devcmd_proxy_by_index) = 0x107cb661 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_alloc_discover) = 0xfb7e4707 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_pdev) = 0xae6ae5c9 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_res) = 0xd910c86b is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_res_bar) = 0x31710a7e is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_res_bus_addr) = 0x7be7a062 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_res_count) = 0x759e4b07 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
        ksym(vnic_dev_get_res_type_len) = 0xd122f0a1 is needed by
```

```
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
      ksym(vnic_dev_unregister) = 0xd99602a1 is needed by
kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64
#
```

Possible Cause

- 1 The enic driver is incorrectly installed.
- 2 The enic driver is not installed.

Solution

Ensure that the correct enic driver has been installed. In addition, make sure of the following:

- Specifically, you must ensure the following: the enic and usnic_verbs drivers must match. If you have a mismatch, you can get the above version errors.
- Specifically, the enic and usnic_verbs that come in the Cisco UCS drivers ISO must be matched together. If you use an enic from one Cisco UCS driver ISO and usnic_verbs from another Cisco UCS driver ISO, it will result in the above version errors.

Problem

Intel IOMMU causes the following warnings:

```
# rpm -ivh kmod-usnic_verbs-1.0.4.318.rhel6u5-1.x86_64.rpm
Preparing... ##### [100%]
 1:kmod-usnic_verbs ##### [100%]
WARNING -
Intel IOMMU does not appear to be enabled - please add kernel parameter
intel_iommu=on to your boot configuration for USNIC driver to function.
#
```

Possible Cause

The Intel IOMMU support is not enabled in the Linux kernel.

Solution

Enable Intel IOMMU driver in the Linux kernel.

Problem

Installing DAT user library can cause the following failed dependencies errors:

```
# rpm -ivh libdaplusnic-2.0.39cisco1.0.0.317-1.el6.x86_64.rpm
error: Failed dependencies:
      dapl is needed by libdaplusnic-2.0.39cisco1.0.0.317-1.el6.x86_64
#
```

Possible Cause

The libdapl is installed without installing the DAT library.

Solution

Install the DAT library.

Problem

When viewing the configuration of Cisco usNIC enabled VICS using `usnic_devinfo`, the command output does not list any usNIC interfaces..

Possible Cause

The RDMA service is not enabled.

Solution

Enable RDMA service using the following commands:

```
# service rdma start
Or
# chkconfig rdma on
```

Problem

Installing the `openmpi-cisco` package causes dependencies error as shown below:

```
# rpm -ivh openmpi-cisco-1.10.1cisco1.0.0.424.rhel7u1-1.x86_64.rpm
error: Failed dependencies:
libfabric-cisco is needed by openmpi-cisco-1.10.1cisco1.0.0.424.rhel7u1-1.x86_64
libfabric.so.1()(64bit) is needed by openmpi-cisco-1.10.1cisco1.0.0.424.rhel7u1-1.x86_64
libfabric.so.1(FABRIC_1.0)(64bit) is needed by
openmpi-cisco-1.10.1cisco1.0.0.424.rhel7u1-1.x86_64
```

Possible Cause

`libfabric-cisco` package is not installed or not present.

Solution

Install `libfabric-cisco` package before installing `openmpi-cisco` package

Problem

Removing the `libfabric-cisco` package causes the following failed dependencies errors:

```
# rpm -e libfabric-cisco-1.1.0cisco1.0.0.385.rhel7u1-1.e17.x86_64
error: Failed dependencies:
libfabric-cisco is needed by (installed) openmpi-cisco-1.10.0cisco1.0.0.385.rhel7u1-1.x86_64
libfabric.so.1()(64bit) is needed by (installed)
openmpi-cisco-1.10.0cisco1.0.0.385.rhel7u1-1.x86_64
libfabric.so.1(FABRIC_1.0)(64bit) is needed by (installed)
openmpi-cisco-1.10.0cisco1.0.0.385.rhel7u1-1.x86_64
```

Possible Cause

Cannot remove `libfabric` rpm.

Solution

Remove the `openmpi-cisco` package before you remove `libfabric` package

