

HP XC System Software

Incorporating External NICs HowTo

Version 3.1

Published: January 2008
Edition: 1



© Copyright 2007, 2008 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. UNIX is a registered trademark of The Open Group.

Linux is a U.S. registered trademark of Linus Torvalds.

other attribution goes here

Table of Contents

Overview.....	7
1 Minimum Requirements.....	7
2 Impact.....	7
3 Before You Begin.....	7
3.1 Gathering Node-Specific Information.....	8
3.2 Determining NIC-Specific Information.....	8
3.3 Gathering Networking Information.....	11
3.4 Understanding the platform_vars.ini File.....	11
3.5 Understanding the device_config Command.....	13
4 Applying the HowTo.....	14
4.1 Consolidating Data for the HowTo.....	14
4.2 Edit the platform_vars.ini Configuration File.....	14
4.2.1 Verifying Your Edited File.....	15
4.3 Update the Database for the External Network Card.....	16
4.4 Update the Firewall Custom Configuration.....	17
4.4.1 Node-Specific Format for the -i Option.....	18
4.4.2 IPv6 Configuration.....	18
4.4.3 Verifying the Updated CMDB.....	18
4.5 Reconfiguring the Nodes.....	20
5 Verifying Success.....	20
5.1 Verifying the Ethernet Port.....	20
5.2 Verifying the Ethernet Device.....	20
6 Troubleshooting.....	21
7 On Successful Completion.....	21
8 HP Encourages Your Comments.....	21

List of Tables

1	Minimum Requirements.....	7
2	Affect on Existing HP XC Systems.....	7
3	Modelmap Values.....	12
4	Data Worksheet.....	14

Overview

This document describes the procedures for incorporating more than one external networking interface card (NIC) in a node of an HP XC system. The `cluster_config` utility configures the first external physical Ethernet port.

HP recommends that the ordering of NICs remain consistent for all nodes of the same model type. That is, for nodes of the same model type, the NICs in the same physical location should be configured consistently as part of the same logical network because they typically share the same gateway, netmask, and other settings.

Incorporating NICs involves:

- Gathering configuration and networking data
- Editing configuration files
- Updating the configuration and management database (CMDB)
- Optionally updating the HP XC IP firewall
- Verifying success

Perform this HowTo after installing the physical Network Interface card (NIC) according to the manufacturer's instructions.

1 Minimum Requirements

Table 1 lists the minimum requirements to help you determine if this HowTo is right for you.

Table 1 Minimum Requirements

Component	Minimum Requirement
HP XC System Software Version	HP XC System Software Version 3.1 and higher
Supported Cluster Platforms	<ul style="list-style-type: none">• HP CP3000 (Intel Xeon)• HP CP4000 (AMD Opteron)• HP CP6000 (Intel Itanium 2)
Firmware Requirements	The master firmware tables are available at the following Web site: http://www.docs.hp.com/en/highperfcomp.html .
Knowledge and Experience Level	You should be an experienced Linux system administrator or network administrator to perform the tasks in this HowTo.

2 Impact

Table 2 lists the effects of this HowTo on existing HP XC systems.

Table 2 Affect on Existing HP XC Systems

Affected Function	Description
System availability	The node or nodes that receive the additional network interface cards will be unavailable for a period of time while they are reimaged.
Hardware configuration	Recabling the NICs is required for their physical installation. If you decide to reposition a NIC on the node, additional recabling will be required.

3 Before You Begin

Before you apply this procedure to incorporate an additional NIC, HP recommends that you understand the following information and perform preliminary tasks in this section.

To perform the necessary tasks for this HowTo, you need:

- To gather node-specific information
- To determine NIC-specific information
- To specify networking information
- To understand the platform_vars.ini Configuration File
- To understand how to use the device_config Command

3.1 Gathering Node-Specific Information

The following node-specific information is required to complete the tasks in this HowTo:

- The node name. That is, how the node is known to the HP XC system.

If the nodes have been discovered, entering the `nodename` command on a node returns its node name.

- The model type of the node

Model numbers are displayed on the front of the node.

If the nodes of the HP XC system have been discovered, you can use the `shownode config` command to determine the model numbers for the nodes in an HP XC system, as follows:

```
# shownode config | less
...
nodes:
...
  n1:
    ...
    server_type:          server rx2600
    ...
```



IMPORTANT: The instructions to incorporate an NIC in an Integrity model type rx8620 server differs slightly from other model types.

- The interconnect switch used

The configuration file you will update takes into consideration if Gigabit Ethernet or another interconnect switch, such as Quadrics QSnet^{II}, is used.

- The firewall configuration

By default, the firewall is configured automatically to open all external physical Ethernet ports the same way: that is, virtual ports 22 (`ssh`) and 443 (`https`) are open by default. The firewall prototype files are capable of configuring up to five external physical Ethernet ports per node automatically as shipped; you can easily extend these files for additional external physical Ethernet ports.

It is possible, however, that you need to custom configure the firewall for the extra external network connections. In this version of the HP XC system software, services may not have been modified to use the external ports automatically and you may need to configure the firewalls appropriately for them.

Complete the corresponding portions of Table 4 (page 14) with the information from this section.

3.2 Determining NIC-Specific Information

For most model types, the following Ethernet interface data for the NIC is required:

- the PCI bus ID



NOTE: The PCI bus ID does not apply to the model type rx8620 server.

- the Ethernet device port
- the Ethernet device's MAC address

On all model types except the rx8620, use the `lshw` command on the target node to return the data for the Ethernet interfaces:



IMPORTANT: Do not use the `lshw` command to determine the Ethernet data for model type rx8620 nodes. This action could cause those nodes to crash. See “Considerations for the Model Type rx8620” for information on how to gather the data for the Ethernet interfaces.

```
# /usr/sbin/lshw -class network
...
description: Ethernet interface
product: NetXtreme BCM5721 Gigabit Ethernet PCI Express
vendor: Broadcom Corporation
physical id: 0
bus info: pci@02:00.0           1
logical name: eth0             2
version: 11
serial: 00:00:00:00:00:00      3
...
```

- 1 This line identifies the PCI bus ID. The PCI bus ID follows the text “pci@”.



NOTE: PCI bus IDs have the following format: `xx:yy.z`

If two PCI bus IDs are identical except for the `z` parameter, those buses are on the same board. For example:

```
02:00.0
05:00.0
05:00.1
```

The first bus is on a separate board from the other two, which are usually on the same board.

- 2 This line identifies the Ethernet device port.
3 This line returns the MAC address of the Ethernet device.

If the node is up, run this command to gather the communication data for all the Ethernet devices on this node. For example:

```
# /usr/sbin/lshw -class network
...
description: Ethernet interface
product: NetXtreme BCM5721 Gigabit Ethernet PCI Express
vendor: Broadcom Corporation
physical id: 0
bus info: pci@02:00.0
logical name: eth0
version: 11
serial: 00:00:00:00:00:00

description: Ethernet interface
product: NetXtreme BCM5721 Gigabit Ethernet PCI Express
vendor: Broadcom Corporation
physical id: 0
bus info: pci@03:00.0
logical name: eth1
version: 11
serial: 00:00:00:00:00:01
```

```

description: Ethernet interface
product: NetXtreme BCM5703 Gigabit Ethernet
vendor: Broadcom Corporation
physical id: 1
bus info: pci@05:01.0
logical name: eth2
version: 10
serial: 00:00:00:00:00:02
...

```

The *dd* variable was used in this example to denote the last portion of the MAC address.

Record the information, as shown here:

Ethernet Device	PCI Bus ID	MAC Address
eth0	02:00.0	00:00:00:00:00:00
eth1	03:00.0	00:00:00:00:00:01
eth2	05:01.0	00:00:00:00:00:02

Complete the corresponding portions of Table 4 (page 14) with the information from this section.

Considerations for the Model Type rx8620

You cannot use PCI bus IDs to configure an additional NIC for the model type rx8620; instead you must use the Ethernet device name. Use the `ifconfig -a` command to determine the MAC address for each Ethernet device:

```

# ifconfig -a | grep eth
eth0      Link encap:Ethernet  HWaddr 00:00:00:00:00:10
eth1      Link encap:Ethernet  HWaddr 00:00:00:00:00:20
eth2      Link encap:Ethernet  HWaddr 00:00:00:00:00:30
...

```

The value following `HWaddr` in this line is the MAC address for the corresponding Ethernet device. If you are unsure of which Ethernet device is which, use the following procedure:

1. Determine the number of Ethernet devices from the `ifconfig -a` command. In this example, there are three: `eth0`, `eth1`, and `eth2`.
2. Disconnect one of the Ethernet devices from the node and take note which one you disconnected.
3. Enter `ethtool ethn | grep detected` on the command line for each of the Ethernet devices.

In this example, enter:

```

# ethtool eth0 | grep detected
Link detected: no
# ethtool eth1 | grep detected
Link detected: yes
# ethtool eth2 | grep detected
Link detected: yes

```

This indicates that the disconnected Ethernet device is `eth0`.

4. Repeat steps 2 and 3 for the remaining Ethernet devices.

Complete the corresponding portions of Table 4 (page 14) with the information from this section.

3.3 Gathering Networking Information

The following information is required to configure the NIC:

- The external host name for the node
- The external IP address
- Optionally, the external IPv6 address
- The netmask
- The gateway, that is, the system that acts as a gateway to external communications
- Optionally, the largest packet size (MTU).

Complete the corresponding portions of Table 4 (page 14) with the information from this section.

3.4 Understanding the platform_vars.ini File

The `/opt/hptc/config/platform_vars.ini` file on the HP XC head node contains configuration parameters and Ethernet port mappings for the various platforms.

There is a section with the following format for each supported model type:

```
[type modeltype]
type=modeltype
alias=server_type
modelmap = <<EOT
Ethernet    expression  Admin_value  Interconnect_value  External_value
!Ethernet   expression  Admin_value  Interconnect_value  External_value
EOT
```

Depending whether or not Gigabit Ethernet is used for the interconnect switch and depending on the number of Ethernet ports that are detected, the values in the table are assigned to the Administration Port, the Interconnect port, and the External port. These values can be stated as Ethernet device names, PCI bus IDs, and the literal strings `undef` and `offboard`.

The following is an example of an entry for the Model rx2600 in the `[type modeltype]` section of the `platform_vars.ini` file:

```
[type rx2600]           1
type=rx2600            2
alias=server rx2600    3
#interconnect  #of eths  AdminPort  Interconnect  External
#-----  -----  -----  -----  -----
modelmap = <<EOT
Ethernet      ==2      undef      undef      undef      4
Ethernet      ==3      20:02.0    offboard    00:03.0
Ethernet      >3      20:02.0    offboard    offboard
!Ethernet     ==2      20:02.0    undef      00:03.0
!Ethernet     >2      20:02.0    undef      offboard
EOT
```

The portions of the entry are:

- 1 This line begins the stanza for this model type.
- 2 This line specifies the model type.
- 3 This line specifies the alias for this model type; it is used during the discover process.
- 4 The `modelmap` variable is used to define the node connections.

Usually, there is more than one line, each with five columns, which are described in Table 3 (page 12). The first line whose first two columns match the following criteria is the only one

used to determine the values for the Administration port, the Interconnect port, and the External port or ports, which are assigned to the `modelmap` variable.

- Interconnect type.
That is, does the system use Gigabit Ethernet for the Interconnect switch or not?
- Ethernet device count.
That is, how many Ethernet devices are detected?

Table 3 Modelmap Values

Column	Values
<code>interconnect</code>	<p>This column contains one of two values:</p> <p><code>Ethernet</code> Indicating a system that uses Gigabit Ethernet for the system interconnect switch.</p> <p><code>!Ethernet</code> Indicating a system that uses a system interconnect switch other than Gigabit Ethernet, such as Quadrics QSnet^{II} and Myrinet.</p> <p>One of these values matches the Interconnect type for the system. That match helps determine which line is used for the <code>modelmap</code> variable.</p>
<code>#of eths</code>	<p>The entry in this column is an expression that describes the number of physical Ethernet ports (the Ethernet device count).</p> <p>The first expression that matches the Ethernet device count detected by the system helps determine which line is used for the <code>modelmap</code> variable.</p>
<code>AdminPort</code>	<p>The port specified in this column is used to connect the node to the administration network of the HP XC system.</p> <p>Valid values for this column are:</p> <p><code>Bus_ID</code>¹ Indicates the hardware PCI bus ID, for example, <code>20:02.0</code>; this is the most reliable method for designating a physical Ethernet port.</p> <p><code>ethn</code> Indicates an Ethernet device, starting with <code>eth0</code>.</p>
<code>Interconnect</code>	<p>The port specified in this column is used for the system interconnect.</p> <p>Valid values for this column are:</p> <p><code>Bus_ID</code>¹ Indicates the hardware PCI bus ID, for example, <code>20:02.1</code>; this is the most reliable method for designating a physical Ethernet port.</p> <p><code>ethn</code> Indicates an Ethernet device, starting with <code>eth0</code>.</p> <p><code>undef</code> Indicates this value is undefined for interconnect port in this configuration. This is the case when an interconnect other than Gigabit Ethernet is used.</p> <p><code>offboard</code>² Is a placeholder indicating the system chooses from an available added NIC.</p> <p>In this HowTo, you will change this value to a more definitive value.</p>
<code>External</code>	<p>The port specified in this column is used to connect to an external network.</p> <p>Valid values for this column are:</p> <p><code>Bus_ID</code>¹ Indicates the hardware PCI bus ID, for example, <code>00:03.0</code>; this is the most reliable method for designating a physical Ethernet port.</p> <p><code>ethn</code> Indicates an Ethernet device, starting with <code>eth0</code>.</p> <p><code>undef</code> Indicates this value is undefined for external port in this configuration.</p> <p><code>offboard</code>² Is a placeholder indicating that the system chooses from an available added NIC.</p> <p>In this HowTo, you will change this value to a more definitive value.</p>

- 1 The entries in the `platform_vars.ini` file for the model type `rx8620` nodes must specify Ethernet devices, not PCI bus IDs.
- 2 This value can be used reliably only when all nodes of that model type have zero or one NIC available after matching up the explicitly specified PCI bus ids to their corresponding logical networks (Admin/Interconnect/External), that is, when the choice of NIC for `offboard` is unambiguous.

The [nodes node] Section

If necessary, you can add one or more `[nodes node]` sections at the end of the `platform_vars.ini` file to configure specific nodes based on the node name instead of the model type. This can be useful as a temporary workaround. The following is an example that configures node `n8`:

```
[nodes n8]
nodes=n8
#interconnect    #of eths    AdminPort    Interconnect    External
#-----
modelmap = <<EOT
Ethernet        ==2        03:00.0      02:00.0         undef
Ethernet        >2        03:00.0      02:00.0         offboard
!Ethernet        >1        03:00.0      undef           02:00.0
EOT
```

This HowTo does not use this section.


3.5 Understanding the device_config Command

The `device_config` command allows you to update the command and management database (CMDB) to incorporate one or more new external Ethernet ports, `external1`, `external2`, and so on. The `cluster_config` utility configures the first external Ethernet port, `external`.

For this HowTo, the `device_config` command has the following format:

```
/opt/hptc/config/sbin/device_config[--dryrun][--type E][--host nodename]
[--extnetwork externalx][--ip IP_address][--name external_hostname]
[--netmask netmask][--gateway gateway][--iftype Ethernet][MAC_address]
```

Where:

<code>nodename</code>	Is the name of the HP XC system node. For example, <code>n19</code> .
<code>externalx</code>	Identifies the external Ethernet port. The first physical external Ethernet port is <code>external</code> , the second port is <code>external1</code> , the next port is <code>external2</code> , and so on.
	NOTE: Do not specify the first physical external Ethernet port, <code>external</code> , with <code>--extnetwork</code> option. This port is the default selection when you do not specify the <code>--extnetwork</code> option.
<code>IP_address</code>	Is the Internet Protocol address for the external connection for the node.
<code>external_hostname</code>	Is the external host name for the node. External systems can access this node directly with this external host name.
<code>netmask</code>	Is the IP network mask for this interface.
<code>gateway</code>	Identifies the system that acts a gateway to external communications.
<code>MAC_address</code>	Is the machine address for the NIC.



NOTE: If you are using IPv6 or MTU, there are additional options that you should specify. Enter the following command for more information:

```
# /opt/hptc/config/sbin/device_config --help
```

4 Applying the HowTo

The following list summarizes the steps to incorporate one or more network interface card for an HP XC system node:

1. “Consolidating Data for the HowTo” (page 14)
2. “Edit the platform_vars.ini Configuration File” (page 14).
3. “Update the Database for the External Network Card” (page 16).
4. “Update the Firewall Custom Configuration” (page 17)
5. “Reconfiguring the Nodes” (page 20)

4.1 Consolidating Data for the HowTo

Enter the data collected in the previous section in Table 4. For more than one node or model type, copy the table as many times as needed.

Table 4 Data Worksheet

Data	Example	Value
Node name	n19	
Model type	rx4640	
Interconnect Switch used	Gigabit Ethernet or Other	
PCI bus ID ¹	03:01.0	
Ethernet device number/port	external1	
Ethernet device's MAC address	00:00:00:00:00:01	
External host name	station1.example.com	
External IP address	192.0.2.10	
External IPv6 address (optional)	FF01:0:0:0:0:0:0:43	
Netmask	255:255:248:0	
Gateway	192.0.2.100	
MTU (optional)		

¹ Not valid for the model type rx8620

4.2 Edit the platform_vars.ini Configuration File

Use the information from the previous section to update the platform_vars.ini configuration file.

1. Log in to the head node as the superuser (root).
2. Change to the /opt/hptc/config directory.
3. Make a backup copy of the platform_vars.ini configuration file, for example:

```
# cp platform_vars.ini platform_vars.ini.BAK
```
4. Use the text editor of your choice to open the platform_vars.ini configuration file.
5. Locate the stanza for the node's model type in the file.

6. Determine which `modelmap` variable assignment line applies to your needs:
 The lines that begin with the string `Ethernet` are provided for systems that use Gigabit Ethernet for the Interconnect switch.
 The lines that begin with the string `!Ethernet` are provided for systems that use an interconnect other than Gigabit Ethernet.
 Determine the line that corresponds to the final number of Ethernet devices for the node.

7. Change the text in the External column accordingly:
 - If the entry in the External column is `offboard`, change that text to the PCI bus ID of the first added NIC.

Before:	After:
<code>offboard</code>	<code>06:01.0</code>

If there is more than one NIC, add a comma character (,) then enter the PCI bus ID of the next added NIC. For example:

Before:	After:
<code>offboard</code>	<code>06:01.0,06:01.1</code>

Repeat as necessary.



IMPORTANT: The list to specify multiple PCI bus IDs must be a comma-separated list without space characters.

- If there already is a PCI bus ID in the External column, append a comma character (,) to that value then enter the PCI bus ID of the first added NIC. For example:

Before:	After:
<code>03:01.0</code>	<code>03:01.0,06:01.0</code>

If there is more than one card, add another comma then enter the PCI bus ID of the next added NIC. For example:

Before:	After:
<code>03:01.0</code>	<code>03:01.0,06:01.0,06:01.1</code>



NOTE: You will need to change the External column of the corresponding line if you change the interconnect switch from Gigabit Ethernet to another interconnect or vice versa.

8. Double-check your entries.
9. Save the file.
10. Exit the text editor.

4.2.1 Verifying Your Edited File

After you have edited the `platform_vars.ini` file, perform the following:

1. From the head node, copy the `platform_vars.ini` to the target node or nodes:

```
# pdcp -w nodelist /opt/hptc/config/platform_vars.ini /opt/hptc/config/
```



TIP: It may be desirable to copy this file to all the nodes.

```
# pdcp -a /opt/hptc/config/platform_vars.ini /opt/hptc/config/
```

2. Log in to the target node.
3. Verify that the `platform_vars.ini` file on the target node contains the changes you made.
4. Run the following `perl` command on the target node. The output helps you determine which device will be selected as the AdminPort, the (first) external device, the interconnect, and any additional external Ethernet devices:

```
perl -I /opt/hptc/perl/lib -e \  
    'use AdapterMap; print join(" ",determineEthMappingFull()),"\n";'
```

```
eth0 eth1 undef eth2 eth3
```

The command output displays the Ethernet devices in the order: Admin, External, Interconnect, External1, External2, and so on. This example shows the following assignments:

Admin	eth0
External	eth1
Interconnect	N/A
External1	eth2
External2	eth3



IMPORTANT: Be sure that the output represents the results you want before proceeding to the next task, “Update the Database for the External Network Card.” Otherwise, repeat the steps in “Edit the `platform_vars.ini` Configuration File” (page 14) until the output accurately represents the results you want.

4.3 Update the Database for the External Network Card

You need to use the `device_config` command on the head node to update the database for an added external NIC.



IMPORTANT: You need to repeat this procedure for each external Ethernet port that you configure.

1. Review the information gathered from Table 4 (page 14). You will need the following information for the `device_config` command:
 - Node name
 - External device/port
 - IP address
 - External host name
 - Netmask
 - Gateway
 - MAC address
2. Enter the `device_config` command with the `--dryrun` option to perform a practice run. For example, the following command is a dry run to update the database for the addition of a NIC for the `external1` port on node `n19` that connects the node for connection as `station1.example.com`:

```
# /opt/hptc/config/sbin/device_config --dryrun \  
    --type E --host n19 \  
    --ip 10.10.10.10 --netmask 255.255.255.0 --gateway 10.10.10.1 --mac 00:00:00:00:00:00
```

```
--extnetwork external1 \  
--ip 192.0.2.2 \  
--name station2.example.com \  
--netmask 255.255.248.0 \  
--gw 192.0.2.100 \  
--iftype Ethernet \  
00:00:00:00:00:02
```



NOTE: The options and arguments are displayed on separate lines for clarity.

The following variables were used in this example:

nn Is the first octet of the IP address and gateway IP address.

dd Is the last portion of the MAC address.

3. Examine the command output.
4. If the command output is correct, run the command again, but without the `--dryrun` option. For example:

```
# /opt/hptc/config/sbin/device_config \  
--type E --host n19 \  
--extnetwork external1 \  
--ip 192.0.2.2 \  
--name station2.example.com \  
--netmask 255.255.248.0 \  
--gw 192.0.2.100 \  
--iftype Ethernet \  
00:00:00:00:00:02
```

4.4 Update the Firewall Custom Configuration

It may be necessary to update the configuration of the IP firewall on the HP XC system. The information in this section may not be needed in all cases.

By default, the virtual IP ports 22 (`ssh`) and 443 (`https`) are open on all physical external Ethernet ports per node. The prototype files for configuring these ports are capable of configuring up to five external Ethernet ports. You can easily edit the `/etc/sysconfig/iptables.proto` file to extend the port configuration.

If a service is not aware of the external physical Ethernet port, it will not be able to communicate through its corresponding virtual ports unless you custom configure the firewall.

As shipped, the firewall prototype file, `/etc/sysconfig/iptables.proto`, contains these lines to configure the firewall:

```
-A RH-Firewall-1-INPUT -i External -p tcp -m tcp --dport 22 -j ACCEPT 1  
-A RH-Firewall-1-INPUT -i External -p tcp -m tcp --dport 443 -j ACCEPT  
-A RH-Firewall-1-INPUT -i External1 -p tcp -m tcp --dport 22 -j ACCEPT 2  
-A RH-Firewall-1-INPUT -i External1 -p tcp -m tcp --dport 443 -j ACCEPT  
-A RH-Firewall-1-INPUT -i External2 -p tcp -m tcp --dport 22 -j ACCEPT 3  
-A RH-Firewall-1-INPUT -i External2 -p tcp -m tcp --dport 443 -j ACCEPT  
-A RH-Firewall-1-INPUT -i External3 -p tcp -m tcp --dport 22 -j ACCEPT 4  
-A RH-Firewall-1-INPUT -i External3 -p tcp -m tcp --dport 443 -j ACCEPT  
-A RH-Firewall-1-INPUT -i External4 -p tcp -m tcp --dport 22 -j ACCEPT 5  
-A RH-Firewall-1-INPUT -i External4 -p tcp -m tcp --dport 443 -j ACCEPT
```

- 1** This line opens virtual port 22 for TCP on the first (non-added) physical external Ethernet port, `External`. The subsequent line performs the same function for virtual port 443.
- 2** This line opens virtual port 22 for TCP on the first additional physical external Ethernet port, `External1`. The subsequent line performs the same function for virtual port 443.
- 3** This line opens virtual port 22 for TCP on the second additional physical external Ethernet port, `External2`. The subsequent line performs the same function for virtual port 443.
- 4** This line opens virtual port 22 for TCP on the third additional physical external Ethernet port, `External3`. The subsequent line performs the same function for virtual port 443.

- 5 This line opens virtual port 22 for TCP on the fourth additional physical external Ethernet port, External4. The subsequent line performs the same function for virtual port 443.

You can modify these lines in the `iptables.proto` file to configure each logical network independently from the others across the HP XC system.

When no Ethernet device serves in the position of a particular network, the line in the `iptables.proto` file is ignored. In effect, the line is dropped from the actual configuration.

See Chapter 11 of the *HP XC System Software Administration Guide* for information on opening IP ports in the firewall.

4.4.1 Node-Specific Format for the `-i` Option

The `iptables.proto` file also has a node-specific format that allows you to control a virtual port for external Ethernet ports only on selected nodes:

```
-i condensed_nodelist[::External[n]]
```

This syntax allows you to open or close ports for a given physical external Ethernet port on the nodes specified by the `condensed_nodelist`. For example:

```
-A RH-Firewall-1-INPUT -i External -p tcp -m tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -i External -p tcp -m tcp --dport 443 -j ACCEPT
-A RH-Firewall-1-INPUT -i External1 -p tcp -m tcp --dport 22 -j ACCEPT 1
-A RH-Firewall-1-INPUT -i External1 -p tcp -m tcp --dport 443 -j ACCEPT 2
-A RH-Firewall-1-INPUT -i n19::External1 -p tcp -m tcp --dport 20,21 -j ACCEPT 3
-A RH-Firewall-1-INPUT -i External2 -p tcp -m tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -i External2 -p tcp -m tcp --dport 443 -j ACCEPT
...
```

- 1 This line opens virtual port 22 for TCP on the first added physical external Ethernet port, External1, on all nodes in the HP XC system.

The text `-i External1` matches all nodes, so virtual port 22 will be open on all nodes with External1 connections.

- 1 This line opens virtual port 443 for TCP on the first added physical external Ethernet port, External1, on all nodes in the HP XC system.

The text `-i External1` matches all nodes, so virtual port 443 will be open on all nodes with External1 connections.

- 3 This line opens the `ftp` virtual ports (20 and 21) on the first added physical external Ethernet port, External1, on node n19.

Additional examples of this syntax are as follows:

```
-A RH-Firewall-1-INPUT -i n23::External1 -p tcp -m tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -i n[5,17]::External1 -p tcp -m tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -i n[9-11,20,33-34]::External1 -p tcp -m tcp --dport 22 -j ACCEPT
```

4.4.2 IPv6 Configuration

If you are using IPv6, you need to configure the `/etc/sysconfig/ip6tables.proto` file. The method for doing so is analogous to configuring the `iptables.proto` file.

4.4.3 Verifying the Updated CMDB

After using the `device_config` command to update the Configuration and Management Database, use the `shownode` command to verify that the database contains the correct data. Enter the following commands to display the data to a file, then use the editor of your choice to examine the file:

```
# shownode config nodes > /tmp/cmdb_outputfile
```

Locate the output that corresponds to the node you changed:

```
n19:
  attributes:
    node_status:          up
```

```

.
.
.
host_names:
  Admin:
    device:                eth0
    gateway:
    hwaddr:                00:00:00:00:00:00
    iftype:                Ethernet
    ifusage:               Admin
    interface_number:
    ipaddr:                172.20.0.16
    ipv6addr:
    mtu:
    name:                  n19
    netmask:               255.224.0.0
    port:                  42
    switch:                172.20.65.1

  External:
    device:                eth1
    gateway:               192.0.2.100
    hwaddr:                00:00:00:00:00:01
    iftype:                Ethernet
    ifusage:               External
    interface_number:
    ipaddr:                192.0.2.1
    ipv6addr:
    mtu:
    name:                  station1.example.com
    netmask:               255.255.252.0

  External1:
    device:                eth2
    gateway:               192.0.2.100
    hwaddr:                00:00:00:00:00:02
    iftype:                Ethernet
    ifusage:               External
    interface_number:
    ipaddr:                192.0.2.2
    ipv6addr:
    mtu:
    name:                  station2.example.com
    netmask:               255.255.248.0

  Interconnect:
    device:                ipoib0
    gateway:
    hwaddr:
    iftype:                Infiniband
    ifusage:               Interconnect
    interface_number:
    ipaddr:                172.22.0.16
    ipv6addr:
    mtu:
    name:                  ic-n19
    netmask:               255.255.0.0
.
.
.

```

4.5 Reconfiguring the Nodes

After editing the `platform_vars.ini` file and updating the database, you need to reconfigure the nodes with the new NICs. Use the following procedure:

1. Log in as superuser (root) on the head node.
2. Copy the following files, and any other files that you may have changed:

```
# pdcp -a -x `hostname` /opt/hptc/config/platform_vars.ini /opt/hptc/config/platform_vars.ini
# pdcp -a -x `hostname` /etc/sysconfig/iptables.proto /etc/sysconfig/iptables.proto
```



NOTE: If you are using IPv6, remember to copy the `ip6tables.proto` file also.

```
# pdcp -a -x `hostname` /etc/sysconfig/ip6tables.proto /etc/sysconfig/ip6tables.proto
```

3. Run the per-node configuration scripts on the node or nodes that you updated in the CMDB:

```
# pdsh -w node "/sbin/service nconfig nconfigure"
# pdsh -w node "/sbin/service nconfig nrestart"
```

Where:

`node` Identifies the node or nodes that were updated.

Alternatively, you can reimage the nodes.

See Chapter 10 of the *HP XC System Software Administration Guide* for information on imaging nodes and distributing software throughout the system and for information on the per-node configuration.

5 Verifying Success

In addition to the verification steps you took after you edited the `platform_vars.ini` file and after you updated the CMDB, you can verify the success of this HowTo by verifying the Ethernet Port settings and the Ethernet device, then connecting to the external network.

5.1 Verifying the Ethernet Port

Use the `ifconfig` command to verify that the Ethernet port for the NIC that you incorporated into the HP XC system is functioning correctly:

```
# ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 00:00:00:00:00:02
          inet addr:192.0.2.2  Bcast:192.0.2.100  Mask:255.255.248.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:10178703  errors:0  dropped:0  overruns:0  frame:0
          TX packets:7442851  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1666416200 (1.5 GiB)  TX bytes:951516823 (907.4 MiB)
          Base address:0x2200  Memory:90300000-90320000
```

For information on the `ifconfig` command, see *ifconfig(8)*.

5.2 Verifying the Ethernet Device

Use the `ethtool` command to verify that the Ethernet device for the NIC is operating correctly:

```
# ethtool eth2
Settings for eth2:
    Supported ports: [ TP ]
    Supported link modes:   10baseT/Half 10baseT/Full
                           100baseT/Half 100baseT/Full
                           1000baseT/Full
    Supports auto-negotiation: Yes
    Advertised link modes:  10baseT/Half 10baseT/Full
                           100baseT/Half 100baseT/Full
                           1000baseT/Full
    Advertised auto-negotiation: Yes
```

```
Speed: 1000Mb/s
Duplex: Full
Port: Twisted Pair
PHYAD: 0
Transceiver: internal
Auto-negotiation: on
Supports Wake-on: umbg
Wake-on: g
Current message level: 0x00000007 (7)
Link detected: yes
```

For information on the `ethtool` command, see *ethtool(8)*.

6 Troubleshooting

Use the standard Linux networking commands to troubleshoot your network connection.

7 On Successful Completion

After you have verified success, be sure to update the golden image to ensure that the changes you made to the configuration files in this HowTo are propagated. See Chapter 10 of the *HP XC System Software Administration Guide* for information on updating the golden image.

8 HP Encourages Your Comments

HP encourages you to provide comments about this document because HP is committed to providing documentation that meets your needs. Send any errors found, suggestions for improvement, or compliments to the following address:

feedback@fc.hp.com

Include the document title, part number, and all other pertinent information in your email.