



Sun™ 10-Port 4X DDR InfiniBand PCIe Network Express Module User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 820-0810-10
March 2007, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.Sun,

Sun Microsystems, the Sun logo, Solaris and Sun Blade are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

The OPEN LOOK and Sun(TM) Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuels relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux Etats - Unis et dans les autres pays.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.Sun,

Sun Microsystems, le logo Sun, Solaris et Sun Blade sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

L'interface d'utilisation graphique OPEN LOOK et Sun(TM) a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

L'utilisation de pieces detachees ou d'unites centrales de remplacement est limitee aux reparations ou a l'echange standard d'unites centrales pour les produits exportes, conformement a la legislation americaine en matiere d'exportation. Sauf autorisation par les autorites des Etats-Unis, l'utilisation d'unites centrales pour proceder a des mises a jour de produits est rigoureusement interdite.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.



Please
Recycle



Adobe PostScript

Contents

Preface vii

1. Introduction to the Sun InfiniBand Network Express Module 1

Product Features 1

Platform and Operating System Support 3

IB NEM Indicators, Buttons, and Ports 4

2. Replacing the Sun InfiniBand NEM 9

Replacing IB NEM Hardware 9

▼ Replace IB NEM in Powered-On Chassis 10

▼ Install IB NEM in Powered-On Chassis 11

Verifying Installation 12

▼ Verify Hardware Installation 13

▼ Verify Installation Using the ILOM Web Interface 13

▼ Verify Installation Using the ILOM CLI 16

▼ Verify Component Status Using the LEDs 18

▼ Verify Installation on Linux 19

▼ Verify Installation on Windows 19

Troubleshooting a Hot-Remove Operation 20

▼ Troubleshoot Using ILOM Web Interface 20

- ▼ Troubleshoot Using ILOM CLI 22
- 3. InfiniBand Software Overview 23**
 - InfiniBand Software for Linux 23
 - OpenFabrics Enterprise Distribution for Linux 25
 - Installing the InfiniBand Drivers on Linux 26
 - ▼ Install IB Drivers From Linux Distribution Source 27
 - ▼ Install the OFED Package 28
 - ▼ Verify Driver Installation on Linux 32
 - InfiniBand Software for Windows 34
 - WinIB for Windows 34
 - ▼ Install the WinIB Package 36
 - ▼ Verify Driver Installation on Windows 41
 - ▼ Verify IPoIB on Windows 42
 - Running OpenSM 43
- 4. Internet Protocol Over InfiniBand 45**
 - Configuring IPoIB on Linux 45
 - ▼ Install IPoIB Driver 46
 - ▼ Change IPoIB Configuration Without Rebooting 47
 - ▼ Change IB Startup Behavior on Linux 48
- 5. Updating IB NEM Firmware 51**
 - Updating IB NEM Firmware for Linux 51
 - ▼ Update IB NEM Firmware for Linux 52
 - Updating IB NEM Firmware for Windows 53
- A. Recommended InfiniBand Cables 55**
 - Passive and Active Copper Cables 55
 - Active Optical Cables 57

Preface

The *Sun 10-Port 4X DDR InfiniBand Network Express Module User's Guide* describes how to install and configure the Sun™ InfiniBand 10-Port Network Express Module (NEM) in a powered-on Sun Blade™ 8000 Series Modular System.

These instructions are designed for enterprise system administrators with experience installing network hardware and software.

How This Document Is Organized

[Chapter 1](#) provides an overview of the Sun InfiniBand (IB) Network Express Module (NEM). This chapter also describes the host platforms and operating systems that support the NEM.

[Chapter 2](#) describes how to replace or install the IB NEM and verify that it has been installed correctly. It also describes how to remove the NEM.

[Chapter 3](#) provides an overview and installation instructions for the InfiniBand software stack for the Red Hat Enterprise Linux, SUSE Linux, and Windows Server 2003 operating systems.

[Chapter 4](#) describes configuration aspects of running the Internet Protocol over InfiniBand (IPoIB).

[Chapter 5](#) provides information on updating the IB NEM firmware on Red Hat Enterprise Linux, SUSE Linux, and Windows Server 2003.

Using UNIX Commands

This document might not contain information about basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at:

<http://docs.sun.com>

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Typographic Conventions

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>% You have mail.</code>
AaBbCc123	What you type, when contrasted with on-screen computer output	<code>% su</code> Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

* The settings on your browser might differ from these settings.

Related Documentation

The documents listed as online are available at:

<http://www.sun.com/documentation/>

Title	Part Number	Format	Location
<i>Sun 10-Port 4X DDR InfiniBand Network Express Module User's Guide</i>	820-0810	PDF	Online
Sun Blade 8000 Series Installation Guide	819-5647	Printed PDF	Online
Sun Blade 8000 Series Online Information System	819-5846	HTML	Online
Sun Blade 8000 Series Product Notes	819-5651	Printed PDF	Online

Documentation, Support, and Training

Sun Function	URL
Documentation	http://www.sun.com/documentation/
Support	http://www.sun.com/support/
Training	http://www.sun.com/training/

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback:

Sun 10-Port 4X DDR InfiniBand Network Express Module User's Guide, Sun part number: 820-0810-10.

Introduction to the Sun InfiniBand Network Express Module

This chapter provides an overview of the Sun 10-Port 4X DDR InfiniBand (IB) PCI Express (PCIe) Network Express Module (NEM). It also lists the host platforms and operating systems that support the IB NEM.

This chapter includes the following sections:

- [“Product Features” on page 1](#)
- [“Platform and Operating System Support” on page 3](#)
- [“IB NEM Indicators, Buttons, and Ports” on page 4](#)

You can order additional Sun 10-Port IB NEMs from Sun Microsystems using the following Marketing part number: X1289A-Z.

Product Features

The Sun IB NEM is a 10-port Network Express Module (NEM) that provides one InfiniBand DDR 4X connection to each Sun Blade Server Module. Each 4X DDR port supports a single x8 PCI Express (PCIe) connection to its Server Module, providing full-duplex data transfers of up to 16 Gbps.

The NEM supports 10 independent InfiniBand Host Channel Adapter (HCA) ports operating at up to 20-Gbps (DDR) speed. It is also backward compatible with 10-Gbps (SDR) devices.

To connect to the external InfiniBand network, standard InfiniBand connectors (microGiGaCN™) are used with either passive copper cables or, optionally, active cables (optical or electrical). [FIGURE 1-1](#) shows the IB NEM back panel.

FIGURE 1-1 IB NEM Back Panel



TABLE 1-1 lists additional features of the Sun IB NEM.

TABLE 1-1 Sun IB NEM Features

Feature	Description
Hot-plug operations supported	Enables system administrators to easily add or remove IB NEMs, as needed, without powering down the system
IB transfer rate (maximum)	5 Gbps (DDR)
InfiniBand Trade Association (IBTA) interoperability	Version 1.2
IB uplink interface	10 InfiniBand 4X copper cable connectors (microGiGaCN) with active cable support
Host interface	10 PCIe x8 links
PCIe transfer rate (maximum)	2.5 Gbps (PCIe Generation One)
IB interface chip	Mellanox InfiniHost™ III Lx (MT25204) HCA
Onboard memory	One 2-MB flash ROM per HCA, 20 MB total field-programmable
Maximum IB cable length	10 meters, using passive copper cables 25–30 meters, using active copper cables Up to 100 meters, using optical cables (DDR speed)
LED indicators	<ul style="list-style-type: none"> • Two LEDs per IB connector (yellow, green) as link physical and logical status indicators • Attention push button, Locate push button and LED (white), OK LED (green), Service Action Required LED (amber), Ready-to-Remove LED (blue)
Form factor	Sun proprietary form factor
RoHS	Fully compliant
Power consumption	Approximately 55 watts total
Protection	Protected from electrostatic discharge (ESD) and handling damage

Platform and Operating System Support

This section provides information about selected platforms that are compatible with the heterogeneous InfiniBand network design.

TABLE 1-2 lists operating system support for the IB NEM on the Sun Blade 8000 Series Modular System.

TABLE 1-2 Operating System Support

Operating System	Version Details
Red Hat Enterprise Linux	Advanced Server v.4, Update 3 (RHEL AS 4-U3) or later (32-bit and 64-bit)
SUSE Linux	Enterprise Server 9 Service Pack 3 (SLES9 SP3) or later (64-bit) Enterprise Server 10 (SLES10) (64-bit)
Microsoft Windows	Microsoft Windows Server 2003 Standard Edition with Service Pack 1 Microsoft Windows Server 2003 Enterprise Edition with Service Pack 1 Microsoft Windows Server 2003 Standard Edition R2 Microsoft Windows Server 2003 Enterprise Edition R2 Microsoft Windows Server 2003 Standard Edition x64 Microsoft Windows Server 2003 Enterprise Edition x64 Microsoft Windows Server 2003 Standard Edition x64 R2 Microsoft Windows Server 2003 Enterprise Edition x64 R2

Note – Currently, Solaris™ Operating System support for the IB NEM is not available. For the latest support information, see the *Sun Blade 8000 Series Product Notes* (Sun Part number: 819-5651) or the following Sun web site:

<http://www.sun.com/documentation/>

TABLE 1-3 provides additional details on the operating systems that provide support for the IB NEM.

TABLE 1-3 Supported Operating Systems and Drivers Required

Operating Systems	Driver Required
RHEL AS 4-U3 or later	Installed with OS
SLES9 SP3 or later, SLES10	Installed with OS
Microsoft Windows 32-bit	Winib_X86-1_3_1347.msi
Microsoft Windows 64-bit	Winib_X64-1_3_1347.msi

For more information about the InfiniBand software stack and related topics, see [Chapter 3](#).

IB NEM Indicators, Buttons, and Ports

As shown in [FIGURE 1-2](#), the Sun IB NEM provides a standard set of status indicators, buttons, and ports on the back panel. [TABLE 1-4](#) lists and describes the components indicated by the numbers.

FIGURE 1-2 Sun IB NEM Back Panel Indicators, Buttons, and Ports




TABLE 1-4 Sun IB NEM Back Panel Components and Functions

Number	Component Name	Color	Description
1	Ejector lever		To remove the NEM, you open this ejector lever.
2	Ready-to-Remove indicator	Blue	Provides the following indications: <ul style="list-style-type: none"> • Steady On – Lights when it is safe to remove the associated NEM from the NEM slot in the chassis. • Off – The NEM is not ready for removal. This indicator is normally lit from the ILOM web interface through the Chassis Monitoring Module (CMM) or by pressing the Attention button.
3	Service Action Required indicator	Amber	Provides the following indications: <ul style="list-style-type: none"> • Steady On – Lights when there is a fault associated with the NEM. • Off – The NEM has no fault condition.

TABLE 1-4 Sun IB NEM Back Panel Components and Functions *(Continued)*

Number	Component Name	Color	Description
4	OK indicator	Green	<p>Displays the different states of module initialization.</p> <p>Provides the following indications:</p> <ul style="list-style-type: none">• Steady On – Lights steadily when NEM is operating normally.• Slow Blink – Blinks slowly when NEM is transitioning from one state to the next.• Standby Blink – Blinks more slowly when NEM has been inserted and NEM is waiting for the Attention button to be pressed.• Off – NEM has no power or one of the other LEDs is lit. <p>The primary purpose of this OK indicator is to provide a visual cue to the operator as to what will happen when using the Attention button.</p>
5	Locate button and indicator	White	<p>The Locate indicator is a bright white button with an associated indicator that blinks (4 Hz) when initiated (remotely) from the ILOM web interface or when the Locate button is pressed (locally).</p> <p>The Locate indicator enables a system operator to easily locate a NEM in a system within a large data center. Once activated, this indicator will time-out after 30 minutes.</p> <p>When pressed, the Locate button lights the Locate indicator.</p> <p>The Locate indicator provides these indications:</p> <ul style="list-style-type: none">• Fast blink – Identifies a specific NEM in the chassis. Lights when the LED is initiated from the web interface remotely or from a press of the Locate button locally.• Off – NEM locator function has not been selected.
6	Attention button		<p>When pressed, the Attention button enables you to activate (if inserting) or prepare to remove (if removing) a NEM during a hot-plug operation.</p>
7	Connectors		<p>InfiniBand 4X copper cable connectors (microGiGaCN) with media detect circuit for optional fiber media adapter.</p>

Each InfiniBand port has two LED indicators, as shown in [FIGURE 1-3](#). [TABLE 1-5](#) lists and describes the LED indicators.

FIGURE 1-3 Sun IB NEM Port LED Indicators




TABLE 1-5 Sun IB NEM Port LED Indicator Descriptions

Number	Color	Description
1	Green	This physical link LED illuminates when the port is electrically active, that is, when a driver is attached and a physical link to a remote switch (or, possibly an HCA) has been established. This indicator is on the left above the cable connector.
2	Yellow	This logical link LED illuminates when the InfiniBand port is in the UP or active state and data can flow through the connection. This indicator is on the right above the cable connector.

Replacing the Sun InfiniBand NEM

This chapter describes how to replace an InfiniBand (IB) NEM in a powered-on Sun Blade 8000 Series Chassis. It also includes instructions to verify that the replacement NEM has been installed correctly.

This chapter contains the following sections:

- [“Replacing IB NEM Hardware” on page 9](#)
- [“Verifying Installation” on page 12](#)
- [“Troubleshooting a Hot-Remove Operation” on page 20](#)



Caution – Damage to the IB NEM can occur as the result of careless handling or electrostatic discharge (ESD). Always handle the NEM with care to avoid damage to electrostatic sensitive components. To minimize the possibility of ESD-related damage, Sun strongly recommends using both a workstation antistatic mat and an ESD wrist strap. You can get an ESD wrist strap from any reputable electronics store or from Sun as part number 250-1007.

You can install the IB NEM in the following Sun Blade 8000 Series Chassis:

- Sun Blade 8000 Chassis
- Sun Blade 8000 P Chassis

Replacing IB NEM Hardware

If an IB NEM fails or if you choose to change the I/O configuration, you will need to replace the NEM. You can replace an IB NEM in a powered-on Sun Blade 8000 Series Chassis using a hot-plug operation.

When you remove a NEM, you must replace it within two minutes to prevent adjacent modules from overheating. If you are removing but not replacing the NEM, you must install a NEM filler panel to meet FCC limits for electromagnetic interference (EMI) and to ensure proper airflow and cooling.

If you encounter a problem replacing the IB NEM, see [“Troubleshooting a Hot-Remove Operation” on page 20](#).

Note – If installing an IB NEM in a Sun Blade 8000 Series Chassis that has not been powered on, see the *Sun Blade 8000 Series Installation Guide* (Sun Part number: 819-5647).

The IB NEMs are customer-replaceable units (CRUs).

▼ Replace IB NEM in Powered-On Chassis

1. Identify which IB NEM to replace.

If the amber Service Action Required LED is lit, this indicates a problem with a specific IB NEM. Otherwise, you can choose any IB NEM to replace if, for example, you want to change the I/O configuration.

2. Prepare the IB NEM for a hot-plug procedure. Use either of these methods:

- Press the Attention button on the IB NEM to initiate the hot-plug removal.

The green OK LED will blink for up to one minute, indicating that the IB NEM is being prepared for removal.

To abort the operation, press the Attention button again within five seconds.

Once the green LED goes dark and the blue LED is illuminated, you can safely remove the NEM.

- Use the ILOM web interface or the command-line interface (CLI) to initiate the hot-plug removal.

If the IB NEM fails the hot-plug preparation and its Ready-to-Remove indicator does not light, see [“Troubleshooting a Hot-Remove Operation” on page 20](#).

3. When the blue Ready-to-Remove LED is lit, physically remove the IB NEM as follows:

- a. Remove all cables from the NEM.
- b. Press the latch on both ejector levers inward at the same time.
- c. Swing out the ejector levers to their fully open position.

d. **Slide the NEM out of its slot.**

Support the weight of the IB NEM with one hand at the bottom of the NEM.

▼ Install IB NEM in Powered-On Chassis

Note – Before installing the replacement IB NEM, locate its GUID (Globally Unique Identifier) printed on the rear of the NEM. Record the GUID and the number of the slot into which you are about to install the NEM and keep this data for future reference.

1. Align the replacement IB NEM with the vacant NEM slot.

Ensure that the port connectors are facing toward you and that the ejector levers are fully open.

2. Align the replacement IB NEM with the chassis guidance system and slide the NEM into its slot until the ejector levers engage and start to close.

Ensure that the NEM engages with the chassis guidance system. Failure to align the NEM correctly can result in damage to the NEM's internal connections to the chassis midplane.

3. Close the levers to secure the IB NEM in its slot. They will click when locked.

Ensure that the back plate on the module mounts flush with the chassis panel opening.

The green OK indicator on the NEM should be in Standby Blink mode.

4. Connect the InfiniBand cable(s) to the IB NEM port connector(s).

5. Ensure that the connector(s) is properly engaged.

The connectors will click when locked.



Caution – Avoid putting unnecessary stress on the connection. Do not bend or twist the cable near the connectors and avoid sharp cable bends of more than 90 degrees.

6. If you have not already done so, connect the other end of the InfiniBand cable(s) to the appropriate port(s) on an InfiniBand switch.

7. Press the Attention button to notify the Sun Blade Server Module(s) (host operating systems) of the NEM.

After you physically install the IB NEM, the Chassis Monitoring Module (CMM) will automatically detect the presence of the NEM. The green OK indicator on the NEM will transition from Standby Blink to Steady On when the NEM is operational.

- To complete installation on a Linux OS, see [“Install the OFED Package” on page 28.](#)
- To complete installation on a Windows OS, see [“Install the WinIB Package” on page 36.](#)

Note – If you are replacing an IB NEM, you will not need to install the InfiniBand software packages. The appropriate software package will have been installed and configured as part of the initial IB NEM installation.

8. Verify that the IB NEM is working properly.

See [“Verify Hardware Installation” on page 13.](#)

- To verify installation on the Linux OS, see [“Verify Installation on Linux” on page 19.](#)
- To verify installation on the Windows OS, see [“Verify Installation on Windows” on page 19.](#)

Verifying Installation

If you have not installed the IB NEM in the chassis and connected it to an operational InfiniBand switch, do so before you attempt to verify the installation. The InfiniBand switch should automatically recognize InfiniBand servers when they are connected to the fabric.

▼ Verify Hardware Installation

1. **Once you have physically installed the IB NEM and ensured that the cables are connected to the IB NEM and switch(es), ensure that an IB subnet manager is running on the connected InfiniBand fabric (network).**

If the green port LED is illuminated, you have successfully completed the hardware installation and you can proceed to verification through the ILOM interface(s). The green LED indicates that the port is enabled, that is, that a physical link to a remote switch (or, possibly an HCA) has been established.

If the port LEDs are **not** illuminated, one possible cause might be that the InfiniBand drivers are not installed. You cannot verify a complete installation on Linux or Windows until you install these drivers.

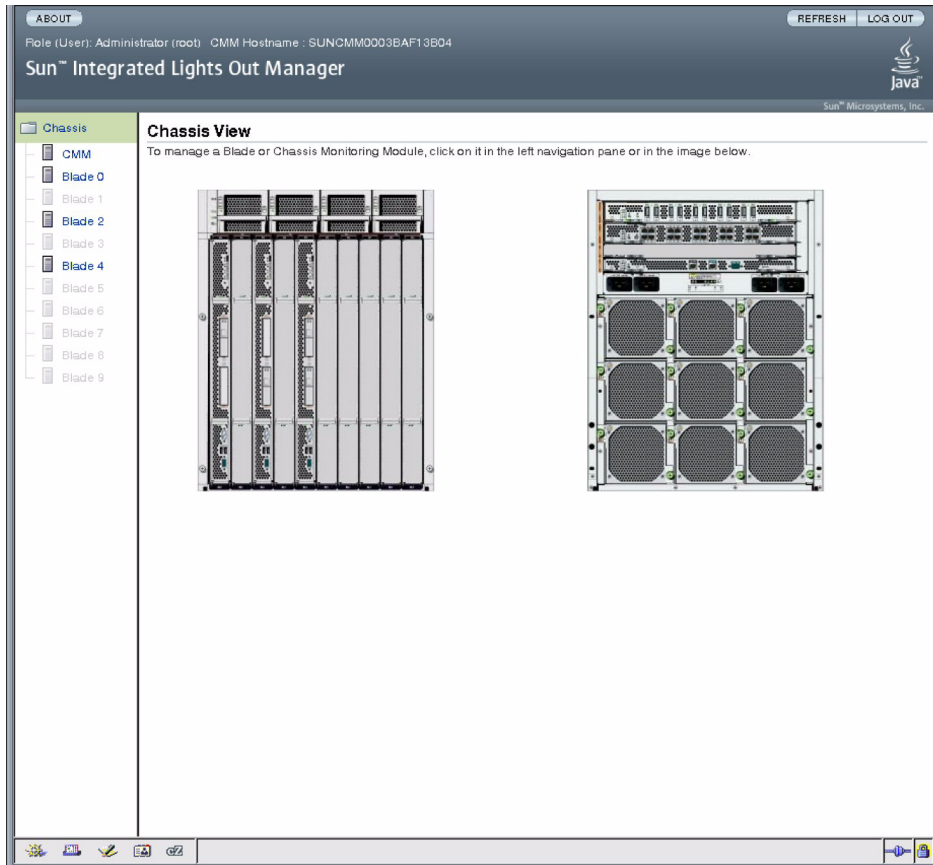
2. **You can now examine hardware status through one of the ILOM (Integrated Lights Out Manager) interfaces. Use one of the following procedures:**
 - [“Verify Installation Using the ILOM Web Interface” on page 13](#)
 - [“Verify Installation Using the ILOM CLI” on page 16](#)

For a description of the possible states of the NEM LEDs, see [“Verify Component Status Using the LEDs” on page 18](#).

▼ Verify Installation Using the ILOM Web Interface

1. **Log in to the ILOM web interface using the IP address of the active CMM.**

The initial page of the ILOM web interface appears, providing visual verification of successful hardware installation. Note the image of the installed IB NEM in the view of the back of the chassis.



- 2. In the left navigation pane, select CMM.**
The ILOM Version Information page appears.
- 3. Select the System Information tab and then select the Components tab.**
The Component Management page appears.

ABOUT REFRESH LOG OUT
 Role (User): Administrator (root) CMM Hostname : SUNCMM0003BAF13B04
 Sun™ Integrated Lights Out Manager Sun™ Microsystems, Inc. Java

Chassis
 CMM
 Blade 0
 Blade 1
 Blade 2
 Blade 4
 Blade 5
 Blade 6
 Blade 7
 Blade 8
 Blade 9

System Information System Monitoring Configuration User Management Remote Control Maintenance
 Versions Session Time-Out Components

Component Management

View component details or prepare to install or remove a component from this page. To modify a component, select the radio button next to that component, then choose an option from the Action drop down list. Components without radio buttons cannot be modified. Choosing the *Prepare to Remove* action shuts down the selected component and lights its blue *Ready to Remove* LED.

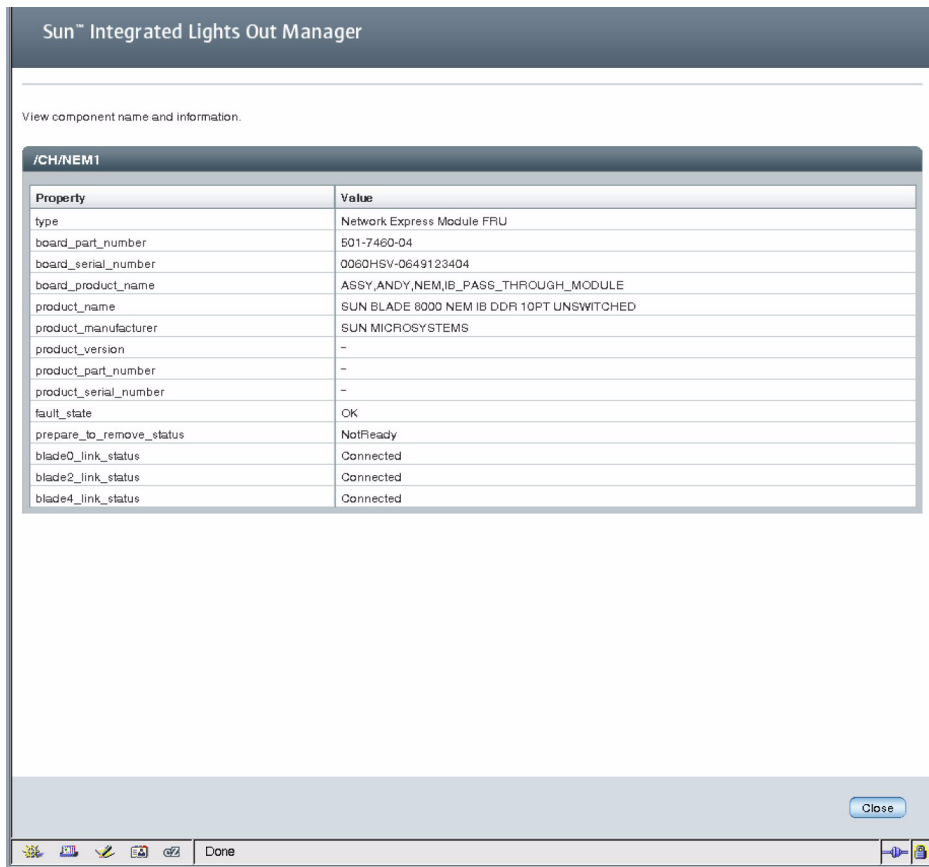
Component Management Status

— Actions —

Component Name	Type	Ready to Remove Status	Fault Status
• /CH	Chassis	-	OK
• /CH/EXTERNAL/AIR_TEMP	Thermal Conditions	-	OK
• /CH/CMM0	Chassis Monitoring Module	Not Ready	OK
• /CH/PS0	Power Supply	Not Ready	OK
• /CH/PS0/EXTERNAL/AC_INPUT	Power Conditions	-	OK
• /CH/PS0/EXTERNAL/AC_INPUT_RANGE	Power Conditions	-	OK
• /CH/PS0/EXTERNAL/AC_INPUT_HALF_LOAD	Power Conditions	-	OK
• /CH/PS1	Power Supply	Not Ready	OK
• /CH/PS1/EXTERNAL/AC_INPUT	Power Conditions	-	OK
• /CH/PS1/EXTERNAL/AC_INPUT_RANGE	Power Conditions	-	OK
• /CH/PS1/EXTERNAL/AC_INPUT_HALF_LOAD	Power Conditions	-	OK
• /CH/PS2	Power Supply	Not Ready	OK
• /CH/PS2/EXTERNAL/AC_INPUT	Power Conditions	-	OK
• /CH/PS2/EXTERNAL/AC_INPUT_RANGE	Power Conditions	-	OK
• /CH/PS2/EXTERNAL/AC_INPUT_HALF_LOAD	Power Conditions	-	OK
• /CH/PS3	Power Supply	Ready (No Power)	OK
• /CH/PS3/EXTERNAL/AC_INPUT	Power Conditions	-	Faulted
• /CH/PS3/EXTERNAL/AC_INPUT_RANGE	Power Conditions	-	OK
• /CH/PS3/EXTERNAL/AC_INPUT_HALF_LOAD	Power Conditions	-	OK
• /CH/RFBD	Rear Fan Board	-	-
• /CH/RFM1	Rear Fan	Not Ready	OK

- Select the IB NEM component name. (You may need to scroll down in the Component Management Status page.)

The ILOM page showing the NEM status details appears.



Note that the NEM `prepare_to_remove` status is `NotReady` and that the `bladen_link_status` is `Connected`, indicating successful hardware installation.

5. If you are physically near the IB NEM, you can examine its LEDs to verify that it has returned the expected feedback.
See “Verify Component Status Using the LEDs” on page 18.

▼ Verify Installation Using the ILOM CLI

1. Log in to the ILOM CLI.
2. To find the IB NEM in your system, enter:


```
> show /CH/
```

3. To verify that the IB NEM is installed, that is, that the Ready To Remove status is Not Ready, enter:

> **show /CH/component**

Where *component* is *NEMn*. The system returns the following message:

```
> show /CH/NEM1
Targets:
SERVICE
OK2RM
LOCATE
OK
ATTN_BTN
LOCATE_BTN
T_AMB
T_CORE

Properties:
type = Network Express Module FRU
board_part_number = 501-7460-04
board_serial_number = 0060HSV-0649123404
board_product_name = ASSY,ANDY,NEM,IB_PASS_THROUGH_MODULE
product_name = SUN BLADE 8000 NEM IB DDR 10PT UNSWITCHED
product_manufacturer = SUN MICROSYSTEMS
product_version = (none)
product_part_number = (none)
product_serial_number = (none)
fault_state = OK
clear_fault_action = (none)
prepare_to_remove_status = NotReady
prepare_to_remove_action = (none)
return_to_service_action = (none)
blade0_link_status = Connected
blade2_link_status = Connected
blade4_link_status = Connected
```

Note that the NEM `prepare_to_remove` status is `NotReady` and that the `bladen_link_status` is `Connected`, indicating successful hardware installation.

4. If you are physically near the IB NEM, you can examine its LEDs to verify that it has returned the expected feedback.

See [“Verify Component Status Using the LEDs”](#) on page 18.

▼ Verify Component Status Using the LEDs

- **Verify the status of the IB NEM using the LEDs.**

TABLE 2-1 lists the possible combinations for the IB NEM LEDs and the status of the NEM indicated by these combinations.

TABLE 2-1 LED Combinations and IB NEM Status

Blue (Top)	Amber (Middle)	Green (Bottom)	IB NEM Status
On	Off	Off	Ready to remove
Off	On	Off	Service attention required
Off	Off	Very Slow Blink	Standby
Off	Off	Slow Blink	Connecting (or disconnecting) links to blades
Off	Off	On	Links connected to blades

To verify successful IB NEM insertion, check the following LEDs and the ILOM interface status:

1. At module insertion, the green OK indicator goes to Standby Blink.
2. When the Attention button is pressed to activate the NEM, the green OK indicator transitions to Slow Blink.
3. When all links to active blades have been made, the green OK indicator transitions to Steady On.
4. In the ILOM interfaces, the Ready to Remove status shows Not Ready.

▼ Verify Installation on Linux

- To determine whether the IB NEM is visible to the Linux OS, enter the `lspci` command.

Output similar to the following appears.

```
> lspci
00:00.0 Memory controller: nVidia Corporation CK804 Memory
Controller (rev a3)
00:01.0 ISA bridge: nVidia Corporation CK804 ISA Bridge (rev a3)
...
01:01.0 SCSI storage controller: LSI Logic / Symbios Logic SAS1064
PCI-X Fusion-MPT SAS (rev 02)
02:00.0 Ethernet controller: Intel
Corporation 82571EB Gigabit Ethernet Controller (rev 06)
02:00.1 Ethernet controller: Intel Corporation 82571EB Gigabit
Ethernet Controller (rev 06)
80:00.0 Memory controller: nVidia Corporation CK804 Memory
Controller (rev a3)
80:01.0 Memory controller: nVidia Corporation CK804 Memory
Controller (rev a3)
80:0e.0 PCI bridge: nVidia Corporation CK804 PCIE Bridge (rev a3)
81:00.0 InfiniBand: Mellanox Technologies MT25204 [InfiniHost III
Lx HCA] (rev 20)
```

The last entry in the sample output (InfiniHost III Lx HCA) verifies the hardware installation and confirms the NEM's availability to the Linux host.

▼ Verify Installation on Windows

You cannot verify installation on the Windows OS until you install the InfiniBand drivers. However, the IB NEM is a Windows Plug and Play device and once the hardware is installed, the Windows OS will detect its presence, and if this is the initial installation, the Windows OS will launch the Found New Hardware Wizard.

- Click **Cancel** to close the Found New Hardware window and exit this instance of the Found New Hardware Wizard.



To install the InfiniBand drivers, launch the Mellanox Microsoft Installer (MSI) for WinIB. See [“Install the WinIB Package”](#) on page 36.

Troubleshooting a Hot-Remove Operation

Because IB NEMs are shared resources, all Sun Blade Server Modules must respond favorably to the PCI hot-remove request. However, a blade might not relinquish the link to a NEM if, for instance, there are busy NFS mounted volumes, file transfers, and so on.

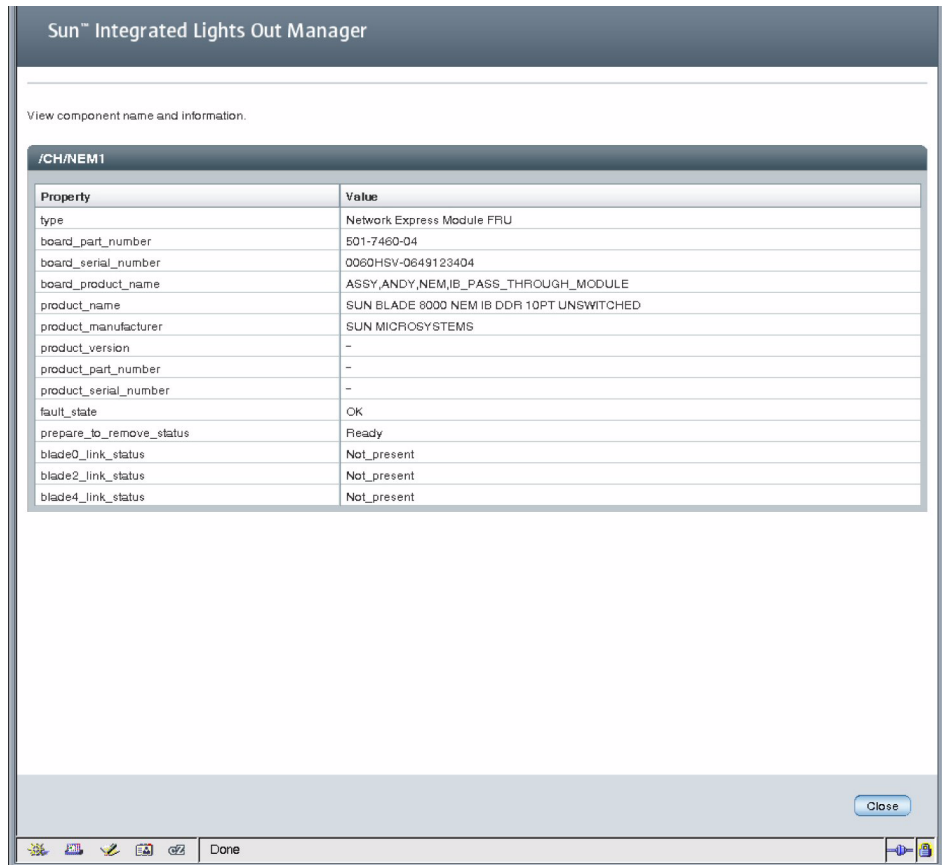
To determine the state of the NEM-to-blade connections, you can use the ILOM web interface or the ILOM command-line interface, as described in the following procedures.

▼ Troubleshoot Using ILOM Web Interface

1. To verify NEM-to-blade connections, log in to the ILOM web interface for the CMM.
2. In the left navigation pane, select CMM.

The ILOM Version Information page appears.

3. **Select the System Information tab and then select the Components tab.**
The Component Management page appears.
4. **Click on the IB NEM component name.**
A page displaying properties and values for the selected IB NEM appears.



As shown, the system responds with `bladen_link_status` entries for each blade (where *n* is the blade module number). Any blade not reporting a `Not_present` status needs intervention from the host OS on that blade. This intervention from the host OS depends entirely on the OS that is active on the blade. Each supported OS has a different method for managing attached devices.

5. **Perform the appropriate host OS procedure for releasing the NEM from the blade.**
6. **Re-execute the steps in “Replace IB NEM in Powered-On Chassis” on page 10.**

▼ Troubleshoot Using ILOM CLI

1. To verify NEM-to-blade connections, log in to the ILOM CLI.
2. Enter the following command, where *n* is the number of the NEM in question.

```
> show /CH/NEMn
Targets:
  SERVICE
  OK2RM
  LOCATE
  OK
  ATTN_BTN
  LOCATE_BTN
  T_AMB
  T_CORE

Properties:
  type = Network Express Module FRU
  board_part_number = 501-7460-04
  board_serial_number = 0060HSV-0649123404
  board_product_name = ASSY,ANDY,NEM,IB_PASS_THROUGH_MODULE
  product_name = SUN BLADE 8000 NEM IB DDR 10PT UNSWITCHED
  product_manufacturer = SUN MICROSYSTEMS
  product_version = (none)
  product_part_number = (none)
  product_serial_number = (none)
  fault_state = OK
  clear_fault_action = (none)
  prepare_to_remove_status = Ready
  prepare_to_remove_action = (none)
  return_to_service_action = (none)
  blade0_link_status = Not_present
  blade2_link_status = Not_present
  blade4_link_status = Not_present
```

As shown, the system responds with `bladen_link_status` entries for each blade (where *n* is the blade module number). Any blade not reporting a `Not_present` status needs intervention from the host OS on that blade. This intervention from the host OS depends entirely on the OS that is active on the blade. Each supported OS has a different method for managing attached devices.

3. Perform the appropriate host OS procedure for releasing the NEM from the blade.
4. Re-execute the steps in [“Replace IB NEM in Powered-On Chassis”](#) on page 10.

InfiniBand Software Overview

InfiniBand is a network architecture that is designed for the large-scale interconnection of computing and I/O nodes through a high-speed switched fabric. To operate InfiniBand on a Sun Blade 8000 Series Modular System, you need an InfiniBand HCA (provided by the IB NEM) and an InfiniBand software stack.

This chapter provides an overview and installation instructions for the InfiniBand software stack for the Linux and Windows operating systems.

Consult the *Sun Blade 8000 Series Product Notes* for the most recent information about supported operating systems, firmware and software updates, and other issues not covered in the main product documentation.

This chapter contains the following sections:

- [“InfiniBand Software for Linux” on page 23](#)
- [“InfiniBand Software for Windows” on page 34](#)
- [“Running OpenSM” on page 43](#)

InfiniBand Software for Linux

If you have installed current releases of Red Hat Enterprise Linux Advanced Server (RHEL AS 4-U3 or later) or SUSE Linux Enterprise Server (SLES9 SP3 or later, SLES10) on a Sun Blade Server Module and you have installed the bundled drivers, you do not need to install or configure additional drivers to support the IB NEM. The bundled drivers will support basic IB NEM operation.

Specifically, RHEL AS 4-U4 contains support in the kernel for HCA hardware produced by Mellanox (mthca driver). The kernel also includes core InfiniBand modules, which provide the interface between the lower-level hardware driver and the upper-layer InfiniBand protocol drivers and provide user space access to InfiniBand hardware.

The kernel also includes the Sockets Direct Protocol (SDP) driver, IP over Infiniband (IPoIB) and the SCSI RDMA Protocol (SRP) driver.

RHEL AS 4-U4 includes the following user space packages:

- `kernel-ib`. Base package that is required to support all other packages. Includes the files necessary to configure the kernel portion of the `openib` stack, create the proper `udev` rules, add the `init` script that allows the kernel modules to be selectively loaded at boot up, and so on.
- `dapl`. RDMA API that supports the DAT 1.2 specification.
- `ibibcm`. InfiniBand Connection Management API.
- `libibcommon`. Common utility functions for the IB diagnostic and management tools.
- `libibmad`. Low-layer IB functions for use by the IB diagnostic and management programs, including MAD, SA, SMP, and other basic IB functions.
- `libibumad`. User MAD library functions that sit on top of the user MAD modules in the kernel. Used by the IB diagnostic and management tools, including OpenSM.
- `libibverbs`. Library that allows user space processes to use InfiniBand "verbs" as described in the InfiniBand Architecture Specification.
- `libibverbs-utils`. Useful subnet and device diagnostic utilities.
- `libmthca`. Device-specific user space driver for Mellanox HCAs (MT23108 InfiniHost and MT25208 InfiniHost III Ex) for use with the `libibverbs` library.
- `libipathverbs`. Device-specific driver for Pathscale HCAs for use with `libibverbs` (only available on x86_64 and ia64 systems).
- `librdmacm`. RDMA Connection Management (`cm`) library.
- `libsdp`. Driver that enables a sockets application to use InfiniBand Sockets Direct Protocol (SDP) instead of TCP transparently and without recompiling the application.
- `openib-diags`. Diagnostic programs and scripts that diagnose the IB subnet.
- `opensm`. Subnet manager software for InfiniBand networks.
- `opensm-libs`. Shared libraries for InfiniBand user space access.
- `perftest`. InfiniBand performance tests.
- `srptools`. In conjunction with the kernel `ib_srp` driver, allows discovery and use of SCSI class devices via the SCSI RDMA Protocol over InfiniBand.
- `mstflint`. Tool to query and update firmware flash memory attached to Mellanox InfiniBand HCAs.

Note – These package names can change, depending on the Linux OS.

The packages selected to support any given configuration will vary. [TABLE 3-1](#) lists the packages considered the absolute minimum needed to support the environment described in this guide.

TABLE 3-1 Required Packages for InfiniBand Support

Package	Command Enabled	Description
kernel-ib	openibd	IB master control script
openib-diags	ibstat	IB utility to display HCAs
openib-diags	ibnetdiscover	IB utility to probe and show the fabric
mstflint	mstflint	Mellanox utility to update HCA FLASHRAM
libibcommon	NA	IB support package
libibmad	NA	IB support package
libibumad	NA	IB support package

If you elected not to install these packages when installing the Linux OS or if you want to upgrade your drivers, you can install these packages at any time from the OS distribution source or by downloading the required files from OpenFabrics.org. For information on both of these procedures, see [“Installing the InfiniBand Drivers on Linux” on page 26](#).

OpenFabrics Enterprise Distribution for Linux

As the popularity of InfiniBand technology increases, the number of Linux distributions and open source organizations producing drivers and tools will increase. For up-to-date information, check with open source organizations and your current vendors.

The OpenFabrics organization is the Open Software solution in the InfiniBand software space and OpenFabrics Enterprise Distribution (OFED) is the InfiniBand suite of software produced by this organization. Various vendors contribute their drivers (and other software components) to OFED.

TABLE 3-2 lists the tested Linux platforms and the corresponding OFED release.

TABLE 3-2 Linux Platforms and OFED Release

Linux Platform	OFED Release
RHEL AS 4-U3 or later	For RHEL AS 4-U3, Sun has tested OFED Release 1.0 of the OpenFabrics stack. RHEL AS 4-U4 includes a built-in version of Release 1.0.
SLES9 SP3 or later, SLES10	Sun has tested OFED Release 1.0 for the SLES10 platform.

OFED contains the following components:

- OpenFabrics core and Upper Layer Protocols (ULPs):
 - HCA drivers
 - Core
 - ULPs, including: IPoIB, SDP, SRP Initiator, iSER Host, RDS and uDAPL
- OpenFabrics utilities:
 - OpenSM (InfiniBand subnet manager)
 - Diagnostic tools
 - Performance tests
- Message Passing Interface (MPI):
 - OSU MPI stack supporting the InfiniBand interface
 - Open MPI stack supporting the InfiniBand interface
 - MPI benchmark tests (OSU BW/LAT, Pallas, Presta)
- Sources of all software modules (under conditions mentioned in the modules' LICENSE files)
- Documentation

Installing the InfiniBand Drivers on Linux

If you did not install the InfiniBand drivers when installing the Linux OS, you can install them at any time from the OS distribution source or by downloading the necessary files from OpenFabrics.org.

To do so, choose one of the following procedures:

- [“Install IB Drivers From Linux Distribution Source” on page 27](#)
- [“Install the OFED Package” on page 28](#)

If you need to determine whether or not the drivers are already installed, see [“Verify Driver Installation on Linux” on page 32](#).

▼ Install IB Drivers From Linux Distribution Source

To install the InfiniBand drivers, you need access to the Red Hat Package Manager (RPM) files. Access to these files is dependent on your individual installation configuration (net boot, CD/DVD boot, .iso files, and so on). When you decide on the appropriate access method and package selection, you can add the packages to the KickStart configuration file (on RHEL) for automatic inclusion in future installations.

Note – On a 32-bit RHEL4 system, all packages have a `.i386.rpm` extension (as shown in the following procedure). On a 64-bit RHEL4 system, all packages have a `.x86_64.rpm` extension instead.

1. Enter the `rpm -ivh` command for each InfiniBand package that you need to install.

Packages must be installed in the following order:

- `libibcommon`
- `libibumad`
- `libibmad`
- `openib-diags`
- `mstflint`
- `perftest`

The following example shows the installation of one package (`libibcommon`) and the resulting dialog on an RHEL AS 4-U4 32-bit system:

```
> rpm -ivh libibcommon-1.0-1.i386.rpm
warning: libibcommon-1.0-1.i386.rpm: V3 DSA signature: NOKEY, key
ID db42a60e
Preparing...      ##### [100%]
1:libibcommon    ##### [100%]
> rpm -ivh libibumad-1.0-1.i386.rpm
.
.
.
```

2. If you are running the CSH or TCSH shell, enter the `rehash` command to rebuild the shell's view of available executables.

3. Enter the `ibstat` command to verify that the OS sees the IB NEM.

```
> ibstat
CA 'mthca0'
  CA type: MT25204
  Number of ports: 1
  Firmware version: 1.1.0
  Hardware version: a0
  Node GUID: 0x001b00000ca72640
  System image GUID: 0x001b00000ca72643
  Port 1
    State: Active
    Physical state: LinkUp
    Rate: 20
    Base lid: 71
    LMC: 0
    SM lid: 2
    Capability mask: 0x02510a68
    Port GUID: 0x001b00000ca72641
```

4. (Optional) You can enter the `ibnetdiscover` command to verify the presence of an operational IB fabric.

For an example of the output of this command, see “[Verify Driver Installation on Linux](#)” on page 32.

5. (Optional) You can check the status of the `ib0` network interface to determine whether the `ib_ipoib` driver is installed.

For details on this step, see “[Install IPoIB Driver](#)” on page 46.

▼ Install the OFED Package

1. On the Sun Blade Server Module, log in as `root` and copy the required files (**OFED1.0.tgz in this example**) from the following location:

<http://www.openfabrics.org/downloads.html>

Note – You will need write access to the files to execute the install script.

2. From `root`, extract the files by entering the following command:
> `tar -zxvf OFED-1.0.tar`
3. From the OFED-1.0 directory, initiate the installation process by entering the following command:
> `./install.sh`

4. When the InfiniBand OFED Distribution Software Installation Menu appears, enter option 2 (Install OFED Software).
5. When the Select OFED Software menu appears, enter option 3 (All packages).
6. When you are asked if you wish to create/install an MPI RPM with gcc, enter **n**.

```
The following compiler(s) on your system can be used to
build/install MPI: gcc
Do you wish to create/install an MPI RPM with gcc? [Y/n]:
```

7. Next, you are asked if you wish to create/install an openmpi RPM with gcc. Again, enter **n**.

```
The following compiler(s) on your system can be used to
build/install openmpi: gcc
Do you wish to create/install an openmpi RPM with gcc? [Y/n]:
```

8. The installation script then lists the OFED packages that it will build. See the following sample output.

```
Following is the list of OFED packages that you have chosen (some
may have been added by the installation program due to package
dependencies):
ib_ipath
ib_ipoib
...
mpitests
ibutils

WARNING: This installation program will remove any previously
installed IB packages on your machine.

Do you want to continue? [Y/n]:
```

9. Enter **Y** to continue.
10. Next, you are prompted to configure InfiniBand IP support. Enter **Y** when asked if you want to include IPoIB configuration files.
Do you want to include IPoIB configuration files (ifcfg-ib*)? [Y/n]:
11. Press Enter to accept the default when prompted to enter a temporary directory for OFED.

```
RPM build process requires a temporary directory.
Please enter the temporary directory [/var/tmp/OFED]:
```

12. Press Enter to accept the default when prompted for the OFED installation directory.

Please enter the OFED installation directory [/usr/local/ofed]:

At this point, the installer begins compiling InfiniBand packages. The process of building packages takes approximately 15–20 minutes.

The system displays output like the following:

```
The MPI_COMPILER_openmpi variable is not defined. Trying the
default compiler: gcc

The following compiler(s) will be used to build the openmpi RPM(s):
gcc

Checking dependencies. Please wait ...

Building InfiniBand Software RPMs. Please wait...

Building openib RPMs. Please wait...
.
.
.
33 packages were built

Build process finished ...
```

13. Installation then begins. See the following message.

```
Removing previous InfiniBand Software installation
Running /bin/rpm -e libibverbs libibverbs-devel libibverbs-
utils...
```

The actual installation takes about one minute.

14. Assuming the IB NEM hardware is installed (and, therefore, an InfiniBand HCA is present), you are prompted to configure InfiniBand IP support. Enter **Y in response to the following prompt:**

Do you want to configure IPoIB interfaces [Y/n]?

The default IPoIB interface configuration is based on DHCP. A special patch for DHCP is required for supporting IPoIB. The patch is available under:

OFED-1.0/docs/dhcp

If you do not have DHCP, you must change this configuration in the following steps.

15. The system next displays the current configuration. When asked if you want to change the configuration as displayed, enter **y**.

```
The current IPOIB configuration for ib0 is:
DEVICE=ib0
BOOTPROTO=dhcp
ONBOOT=yes
Do you want to change this configuration? [y/N]:
```

16. The configuration script guides you through the changes one at a time. See the following as an example.

```
Enter an IP Address:10.0.0.52
Enter the Netmask: 255.255.255.0
Enter the Network:10.0.0.0
Enter the Broadcast Address:10.0.0.255
Start Device On Boot? [Y/n]:Y

Selected configuration:

IPADDR=10.0.0.52
NETMASK=255.255.255.0
NETWORK=10.0.0.0
BROADCAST=10.0.0.255
ONBOOT=yes

Do you want to save the selected configuration? [Y/n]:
```

17. Enter **Y** to save the configuration.

If you have entered a valid IP configuration for `ib0`, you are now properly configured for IPoIB operations.

18. Iterate the InfiniBand configuration over all InfiniBand interfaces.

You will need to enter a valid IP configuration for each network interface.

19. Once all IPoIB interfaces have been configured, you are prompted as follows to configure OpenSM for the blade.

Do you want to configure OpenSM [Y/n]?

20. Enter **n** to complete this part of the installation. You should see a message like the following:

```
Installation finished successfully...
Press Enter to continue...
```

For information on OpenSM, see [“Running OpenSM” on page 43](#).

21. Press Enter.

The InfiniBand OFED Distribution Software Installation Menu is displayed.

22. Enter Q to exit.

The Sun Blade Server Module is configured now to start up the InfiniBand software on reboot (ONBOOT=yes).

If this is not the desired behavior, you can edit the `/etc/infiniband/openib.conf` file, changing ONBOOT to equal no. You can also manually control basic InfiniBand behavior by entering the following command:
`/etc/init.d/openibd { start | stop | status }`

After successful installation, it is recommended that you reboot the Server Module. After reboot, the Server Module should come up as a functional member of the InfiniBand fabric.

▼ Verify Driver Installation on Linux

1. Verify that the Linux software driver is installed and attached to the IB NEM by entering the `openibd status` command.

Note – When using the `openibd` command, enter the entire path as shown in the example.

The following example shows the IB driver installed, running and presenting one IB HCA channel or network device (`ibn`) to the OS. In the example, the Linux network device appears as `ib0`.

```
> /etc/init.d/openibd status
    HCA driver loaded
Configured devices:
ib0
Currently active devices:
ib0
    The following modules are also loaded:
ib_cm
ip_ipoib
.
.
.
```

2. To view details of operational status, enter the `ibstat` command.

The following example shows one operational IB channel into the IB fabric (or network). The `LinkUp` state indicates active participation in an IB fabric. It is present as `lid 69` and it is being managed by `lid 2`.

```
> ibstat
CA 'mthca0'
  CA type: MT25204
  Number of ports: 1
  Firmware version: 1.1.0
  Hardware version: a0
  Node GUID: 0x001b00000ca72620
  System image GUID: 0x001b00000ca72623
  Port 1
    State: Active
    Physical state: LinkUp
    Rate: 20
    Base lid: 69
    LMC: 0
    SM lid: 2
    Capability mask: 0x02510a68
    Port GUID: 0x001b00000ca72621
```

You can also verify that the InfiniBand fabric is operational by entering the `ibnetdiscover` command. The output from this command will list all the nodes, as shown in the following sample output.

```
> ibnetdiscover
#
# Topology file: generated on Thu Jan 11 15:19:59 2007
#
# Max of 4 hops discovered
# Initiated from node 001b00000ca72620 port 001b00000ca72621

vendid=0x8f1
devid=0x5a31
sysimguid=0x8f10400411ef9
switchguid=0x8f10400411ef8

Switch 24 "S-0008f10400411ef8" # Switch port 0 lid 9
[21]      "H-0002c90109761ea0" [2]
[12]      "S-0005ad00000161ba" [5]
[7]       "H-001b00000ca72630" [1]
[6]       "H-001b00000ca72620" [1]
vendid=0x5ad
devid=0xa87c
sysimguid=0x5ad01010161b6
```

```
switchguid=0x5ad00000161ba

Switch 8 "S-0005ad00000161ba" # Switch - U3 port 0 lid 3
[4] " H-0005ad0000011310" [1]
[3] "S-0005ad00000161b6" [1]
[2] "S-0005ad00000161b6" [2]
[1] "S-0005ad00000161b8" [3]
[5] "S-0008f10400411ef8" [12]
.
.
.
vendid=0x2c9
devid=0x6274
sysimgguid=0x1b00000ca72633
caguid=0x1b00000ca72630
Ca 1 "H-001b00000ca72630" # 4x DDR IB 10-Port PCIe Network
Express Module
[1] "S-0008f10400411ef8" [7] # lid 68 lmc 0n
```

Note – The output from `ibnetdiscover` can be lengthy in a large IB fabric.

InfiniBand Software for Windows

If you have installed current releases of Windows Server 2003 on a Sun Blade Server Module, you must also install and configure additional drivers to support the IB NEM. The following section describes that process for Windows.

WinIB for Windows

On the Windows platform, Sun currently redistributes the WinIB Collection version 1.3.0.1347 (32-bit and 64-bit versions) from Mellanox Technologies. WinIB is designed, packaged, and supported to meet the needs of high performance computing (HPC) and enterprise data center (EDC) applications through:

- Collaboration with OpenIB
 - WinIB includes qualified and released versions of OpenIB InfiniBand stack components.
- Support for multiple industry-standard interfaces that enable off-the-shelf applications to work with InfiniBand

- Support for Microsoft Windows Server 2003 and Windows Compute Cluster Server (CCS) 2003

TABLE 3-3 provides details about the InfiniBand software components in the WinIB Collection version 1.3.0.1347.

TABLE 3-3 WinIB Software Components

Component	Description	Version
IBAL (InfiniBand Access Layer) driver, Access Layer, Upper Layer Protocols	Core – IB Access Layer, lower-level drivers.	1.3.0.1347
IPoIB protocol	IP communications over IB fabric.	1.3.0.1347
OpenSM for Windows	IB-compliant subnet manager (SM) and administrator running on top of OpenIB.	1.8.1
Performance test (<code>perf_main</code>)	Application reporting latency and bandwidth HCA performance.	
Documentation	IBAL API HTML documentation, license file, code examples.	

Hardware and software requirements for the Sun Blade Server Module are as follows:

- Installed IB NEM
- Windows Server 2003 (32-bit or 64-bit)
- Administrator privileges on your Server Module
- Disk space for installation: 10 Mbytes

After you have met these requirements, download and install one of the WinIB packages for Windows from the Sun Blade 8000 Series Resource CD.

TABLE 3-4 summarizes Windows platforms and the required InfiniBand drivers. For additional platform details, see TABLE 1-2.

TABLE 3-4 Windows Server 2003 Platforms and Required Drivers

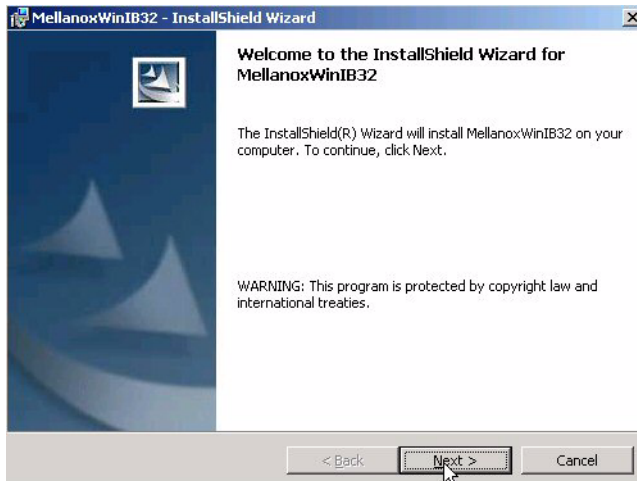
Windows Platform	Driver Required
All supported 32-bit versions	<code>Winib_X86-1_3_1347.msi</code>
All supported 64-bit versions	<code>Winib_X64-1_3_1347.msi</code>

▼ Install the WinIB Package


1. Download the appropriate Microsoft Installer (either x86 or x64) from the Sun Resource CD. Launch the Microsoft Installer.
2. When the Found New Hardware Wizard appears, select No, not this time and click Next to continue.



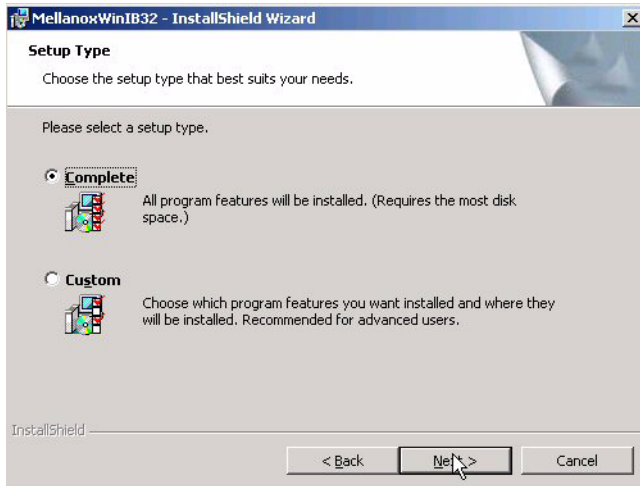
3. In the Welcome screen, click Next to continue.




4. The License Agreement screen appears. Read the license agreement carefully. If a printer has been configured on the blade, you can click Print to send the agreement to the default printer. You must accept the license agreement to continue. After selecting the acceptance radio button, click Next to continue.
5. The Customer Information screen is displayed. Enter your User Name and Company information and click Next to continue.
6. The Destination Folder screen is displayed. Confirm the folder location or change by clicking the Change button and following the subsequent instructions. Click Next to confirm and continue.



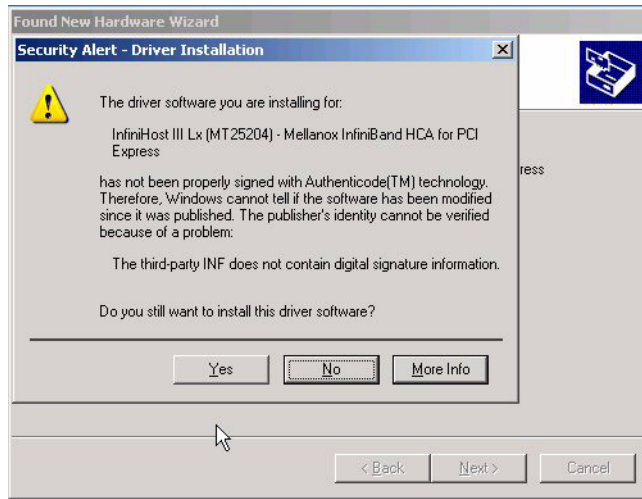
7. The Setup Type screen appears. Select the Complete radio button and click Next to continue.



8. The Ready to Install screen appears. Click Back if you need to review or change any of your selections. When ready, click Install to begin installation.



9. The setup process begins. A cmd window opens in the background and minimizes to the taskbar. Setup pauses at driver installation because the driver is not digitally signed and a Security Alert is displayed on top of the Found New Hardware screen. Click Yes to acknowledge the warning and continue with the installation.



10. Setup displays another Security Alert when installing the InfiniBand Fabric driver because that driver is also unsigned. Click Yes to acknowledge the warning and continue. A long pause follows (approximately 40 seconds). Be patient and take NO action.



11. After the pause, the Found New Hardware Wizard appears again. Do not close this screen



Caution – Do not close this screen or respond to it in any way except to move it aside.

12. The Hardware Installation alert for the OpenIB IPoIB Adapter is displayed. Click Continue Anyway.



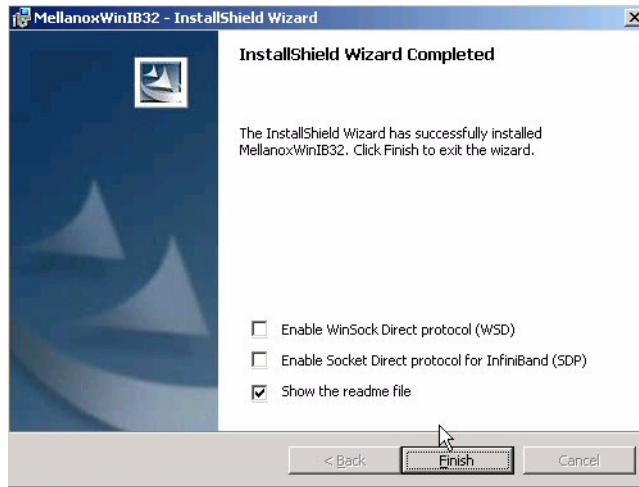
If you do not see this alert displayed, it might be because it is hidden behind the InstallShield screen.



Caution – Do not attempt to close or minimize any of the visible screens. If the Hardware Installation alert is not visible, use the mouse to drag the InstallShield screen aside to reveal it.

Once you click Continue Anyway, all screens will close and the InstallShield Wizard Completed screen is displayed.

13. In the InstallShield Wizard Completed screen, select the desired protocol depending on the requirements of your application (either WSD or SDP for InfiniBand). Leave Show the readme file selected and click Finish to exit the wizard.

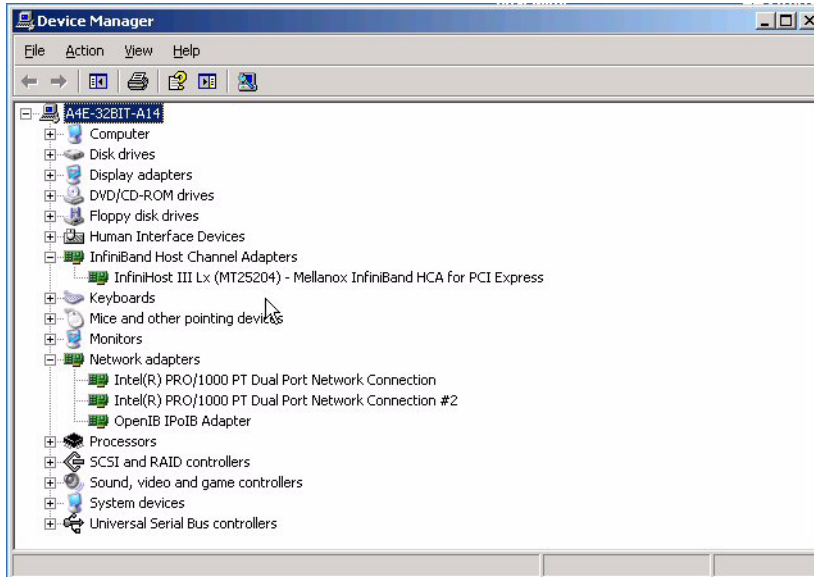


14. **The README.txt file is displayed. Read this file for important last-minute updates.**

It is recommended that you reboot the Sun Blade Server Module at the end of the installation.

▼ Verify Driver Installation on Windows

1. **On Windows, driver installation is verified in the Device Manager. To open the Windows Device Manager, click Start and then click Control Panel.**
2. **Double-click System.**
The System Properties screen opens.
3. **Select the Hardware Tab and click Device Manager.**
The Device Manager screen opens, displaying a new hardware category, InfiniBand Host Channel Adapters. Under this category, InfiniHost III Lx (MT25204) - Mellanox InfiniBand HCA for PCI Express appears. Under the category Network adapters, OpenIB IPoIB Adapter appears.



You can also verify that the TCP/IP service is working (active). Use the following procedure to determine whether the InfiniBand driver is running on the node.

▼ Verify IPoIB on Windows

- **To verify that the IB NEM is installed and that the lower-level InfiniBand driver is running on the node, enter the `vstat` command.**

The node should return `1 HCA` found on the first line and the connected port should report a `PORT_ACTIVE` state.

The `vstat` utility, placed under the `\bin` directory, displays the status and capabilities of the IB NEM where the device resides.

```
> vstat
1 HCA found:
    hca_id=InfiniHost_III_Lx0
    pci_location={BUS=0x81,DEV/FUNC=0x00}
    vendor_id=0x02C9
    vendor_part_id=0x6274
    hw_ver=0x20
    fw_ver=1.2.0
    PSID=SUN0010010001
    num_phys_ports=1
        port=1
        port_state=PORT_ACTIVE
        sm_lid=0x0002
        port_lid=0x0046
        port_lmc=0x00
        max_mtu=2048
```

The optional parameter `-v` runs a "verbose" version of this command.

Running OpenSM

In the InfiniBand architecture, a subnet manager (SM) is required for the InfiniBand fabric to function properly. The SM discovers all the nodes on the fabric and assigns the local identifiers (LIDs) in the HCA(s). It also sets up the routing tables in the switches to support routing packets between nodes.

To meet these needs, OFED and WinIB supply OpenSM, an open source subnet manager. OpenSM can initialize and configure the subnet as well as keep the subnet operational when the network topology and nodes change. OpenSM runs as a system daemon on at least one of the host machines in the InfiniBand fabric.

The OpenSM application also contains the subnet administrator (SA), an associated component that acts like a database and can be affected by end node requests. It supports querying as well as event forwarding. Applications send queries to the SA to discover the path records for remote nodes, which are needed to establish connections between endpoints on the fabric.

Note – Two instances of OpenSM running concurrently on the same port will result in a system crash.

For more information on OpenSM, see the ReadMe file for your OS (OFED for Linux and WinIB for Windows).

Internet Protocol Over InfiniBand

This chapter describes configuration aspects of running the Internet Protocol over InfiniBand (IPoIB) and contains the following section:

- [“Configuring IPoIB on Linux” on page 45](#)

Configuring IPoIB on Linux

You might decide to change your IPoIB configuration for a variety of reasons, including the installation of an additional IB NEM.

▼ Install IPoIB Driver

1. Determine whether the IPoIB driver is already installed by entering the `lsmod | grep ib` command.

The output from this command shows all the IB drivers.

In the following sample output, note that the driver, `ib_ipoib`, is not listed.

```
> lsmod | grep ib
ib_sdp                45340 0
rdma_cm              26760 1 ib_sdp
ib_addr              10504 1 rdma_cm
ib_local_sa          14232 1 rdma_cm
findx                 6528 1 ib_local_sa
ib_ipath              70552 0
ipath_core           179652 1 ib_ipath
ib_mthca             139184 0
ib_uverbs            47536 0
ib_umad              19888 0
ib_ucm               21512 0
ib_sa                 18196 2 rdma_cm,ib_local_sa
ib_cm                 39952 2 rdma_cm,ib_ucm
ib_mad               43176 5 ib_local_sa,ib_mthca,ib_umad,ib_sa,ib_cm
ib_core              59520 11
ib_sdp,rdma_cm,ib_local_sa,ib_ipath,ib_mthca,ib_uverbs,ib_umad,
ib_ucm,ib_sa,ib_cm,ib_mad
```

2. To install the IPoIB driver, enter the `modprobe` command:

```
> modprobe ib_ipoib
```

3. Enter the `lsmod | grep ib` command again and note that `ib_ipoib` is now listed.

```
> lsmod | grep ib
ib_ipoib             59800 0
ib_sdp                45340 0
rdma_cm              26760 1 ib_sdp
ib_addr              10504 1 rdma_cm
ib_local_sa          14232 1 rdma_cm
.
.
.
ib_core              59520 11
ib_sdp,rdma_cm,ib_local_sa,ib_ipath,ib_mthca,ib_uverbs,ib_umad,
ib_ucm,ib_sa,ib_cm,ib_mad
```


4. Enter the `ifconfig` command to check for network interface `ib0`.

```
> ifconfig ib0
ib0      Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-
00-00-00-00-00-00
          BROADCAST MULTICAST  MTU:2044  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:128
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
```

Note that network interface `ib0` is present but has no valid IP address.

To assign an address, see [“Change IPoIB Configuration Without Rebooting” on page 47](#).

▼ Change IPoIB Configuration Without Rebooting

1. Start the InfiniBand IP network by entering the `ifconfig` command and assigning a valid IP address for `ib0`.

```
> ifconfig ib0 10.0.0.50/24
ib0      Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-
00-00-00-00-00-00
          inet
addr:10.0.0.50 Bcast: 10.0.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:128
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
```

If `ib0` were unconfigured, it would appear without an IP address as shown in the following output.

```
> ifconfig ib0
ib0      Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-
00-00-00-00-00-00
          BROADCAST MULTICAST  MTU:2044  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:128
          RX bytes:0 (0.0 b)  TX bytes:0 ( 0.0 b)
```

2. Enter the `route` command to verify that the **10.0.0** subnet is available.

The following output shows subnet `10.0.0` present and routed through `ib0`.

```
> route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use If
ace
10.0.0.0 * 255.255.255.0 U 0 0 0 ib
0
10.8.134.0 * 255.255.255.0 0 0 0 et
h0
169.254.0.0 * 255.255.0.0 U 0 0 0 et
h0
default ban3rtr0d0 0.0.0.0 UG 0 0 0 et
h0
```

3. As shown in the following example, you can enter the `ping` command to see another IPoIB node on the **10.0.0** subnet:

```
> ping 10.0.0.79
PING 10.0.0.79 (10.0.0.79) 56(84) bytes of data.
64 bytes from 10.0.0.79: icmp_seq=0 ttl=255 time=1.82 ms
64 bytes from 10.0.0.79: icmp_seq=1 ttl=255 time=0.082 ms
64 bytes from 10.0.0.79: icmp_seq=2 ttl=255 time= 0.062 ms

--- 10.0.0.79 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 0.062/0.655/1.823/0.825 ms, pipe 2
```

At this point, the IPoIB network is active and properly configured without rebooting.

▼ Change IB Startup Behavior on Linux

To change the InfiniBand startup behavior, edit the `openib.conf` configuration file (full path: `/etc/infiniband/openib.conf`).

In the following example, `openib.conf` specifies that whenever the system boots, the InfiniBand services, IPoIB, and the SDP IP service are to start up automatically (`ONBOOT=yes`, `IPOIB_LOAD=yes`, `SDP_LOAD=yes`). However, `openib.conf` specifies that the SRP service is NOT to start up automatically (`SRP_LOAD=no`). You can alter any and all of these parameters.

1. Edit `ipoib.conf`.

```
# Start HCA driver upon boot
ONBOOT=yes
# Load UCM module
UCM_LOAD=no
# Load RDMA_CM module
RDMA_CM_LOAD=no
# Load RDMA_UCM module
RDMA_UCM_LOAD=no
# Load MTHCA
MTHCA_LOAD=yes
# Load IPATH
IPATH_LOAD=yes
# Load IPoIB
IPOIB_LOAD=yes
# Load SDP module
SDP_LOAD=yes
# Load SRP module
SRP_LOAD=no
# Load RDS module
RDS_LOAD=no
```

2. Create (or edit) the `ifcfg-ibn` file to configure an individual network interface.

For each InfiniBand network interface, you will need a corresponding `/etc/sysconfig/network-scripts/` startup file (`ifcfg-ib0`). As an example, the startup file for `ib0`, might look something like the following.

```
more /etc/sysconfig/network-scripts/ifcfg-ib0
DEVICE=ib0
BOOTPROTO=static
IPADDR=10.0.0.50
NETMASK=255.255.255.0
NETWORK=10.0.0.0
BROADCAST=10.0.0.255
ONBOOT=yes
```

The `ONBOOT=yes` parameter indicates that the corresponding IP network interface is to automatically start up when the system boots. Specifying `ONBOOT=no` will "configure" the interface but not start it.

Generally, if you have enabled IPoIB services, the IB stack installation scripts will automatically create the `ifcfg-ibn` configuration files for all IB network interfaces present. If you install an IB NEM **after** you have installed the IB stack, you might need to manually create the `ifcfg-ibn` files for the newly installed network interfaces.

Updating IB NEM Firmware

This chapter provides information on updating the IB NEM firmware on Linux and Windows.

Consult the *Sun Blade 8000 Series Product Notes* for the most recent information about the availability of firmware updates.

This chapter contains the following sections:

- [“Updating IB NEM Firmware for Linux” on page 51](#)
 - [“Updating IB NEM Firmware for Windows” on page 53](#)
-

Updating IB NEM Firmware for Linux

The firmware version on your IB NEM should be ready to use and should not require updating. However, if you want to update the firmware for any reason, you will need to use vendor-specific and (generally) OS-specific firmware updating tools.

For Linux (RHEL AS 4-U3 or later, SLES9 SP3 or later, and SLES10), use the OFED `mstflint` tool to load new IB NEM firmware. The tool, `mstflint`, is available both as part of the bundled software and from the standard OFED stack.

Installed by default, `mstflint` is similar to the Mellanox `flint` tool with the following exception. You must identify the IB NEM in the PCI `bus:dev.fun` format to satisfy the `mstflint` command `-d` device syntax requirement.

▼ Update IB NEM Firmware for Linux

1. Enter the `lspci` command to identify the IB NEM.

In the following example, the IB NEM (Mellanox InfiniHost III Lx HCA) is configured as PCI bus number 81, device 00, function 0 (81:00.0), which is NEM slot 1 in a Sun Blade 8000 P Series Modular System. On your system, you might see a different designation for the NEM.

```
> lspci
...
80:0e.0 PCI bridge: nVidia Corporation CK804 PCIE Bridge (rev a3)
81:00.0 InfiniBand: Mellanox Technologies MT25204 [InfiniHost III
Lx HCA] (rev 20)
```

2. Enter the `mstflint -d` command in the `bus:dev.fun` format.

```
> mstflint -d 81:00.0 -i ibnem.bin burn
Current FW version on flash: 1.1.0
New FW version: 1.2.100

Burn image with the following GUIDs:
Current GUIDs are:
Node: 001b00000ca72680
Port1: 001b00000ca72681
Sys.Image: 01b00000ca72683

Read and verify Invariant Sector - OK
Read and verify PPS/SPS on flash - OK
Burning second FW image without signatures - OK
Restoring second signature -OK

Do you want to continue ? (y/n) [n] : y
```

Note – The GUIDs you will see (Node, Port1, and Sys. Image) during the burn process will differ from those shown in the example.

3. To burn the image, enter `y`.

4. As with any IB NEM FLASHRAM update, you must reset the Server Module (or at least the IB NEM) to load and execute the new firmware image.

5. After resetting the Server Module (or the NEM), enter the `ibstat` command to verify the new firmware version.

```
> ibstat
CA 'mthca0'
CA type: MT25204
  Number of ports: 1
  Firmware version: 1.1.0
  Hardware version: a0
  Node GUID: 0x001b00000ca72600
  System image GUID: 0x001b00000ca72603
  Port 1:
    State: Active
    Physical state: LinkUp
    Rate: 20
    Base lid: 70
    LMC: 0
    SM lid: 2
    Capability mask: 0x02510a68
    Port GUID: 0x001b00000ca72601
```

Updating IB NEM Firmware for Windows

The firmware version on your IB NEM should be ready to use and should not require updating.

Recommended InfiniBand Cables

This appendix provides information to help you select the appropriate cables to support expected performance.

According to the InfiniBand specification, the link must transmit data with a bit error rate (BER) of at least 10^{-12} . The BER can be guaranteed for DDR speed only when recommended cables are used for InfiniBand connections.

This appendix contains the following sections:

- [“Passive and Active Copper Cables” on page 55](#)
- [“Active Optical Cables” on page 57](#)

Passive and Active Copper Cables

The following table contains the list of recommended passive and active copper cables for the IB NEM.

Vendor	Length (Meters)	Gauge	Equalization	Type	Part #
Amphenol	2	28	E	Passive copper	523640005
Amphenol	2	28	N/E	Passive copper	556760005
Amphenol	5	26	N/E	Passive copper	556770011
FoxConn	1	28	N/E	Passive copper	2GFBBBA-09E
FoxConn	2	28	N/E	Passive copper	2GFBBBX-03E
Fujikura	2	28	N/E	Passive copper	FA1-1129-02A NO.001
Fujikura	5	24	N/E	Passive copper	FA1-1129-05A NO.04

Vendor	Length (Meters)	Gauge	Equalization	Type	Part #
Gore	1	30	E	Passive copper	IBN4300-1
Gore	2	28	E	Passive copper	IBN4800-2
Gore	2	30	E	Passive copper	IBN4300-2
Gore	3	28	E	Passive copper	IBN4800-3
Gore	4	28	E	Passive copper	IBN4800-4
Gore	4	30	E	Passive copper	IBN4300-4
Gore	5	28	E	Passive copper	IBN4800-5
Gore	5	30	E	Passive copper	IBN4300-5
Gore	6	26	E	Passive copper	IBN4600-6
Gore	6	28	E	Passive copper	IBN4800-6
Gore	6	30	E	Passive copper	IBN4300-6
Gore	7	26	E	Passive copper	IBN4600-7
Gore	7	28	E	Passive copper	IBN4800-7
Gore	7	30	E	Passive copper	IBN4300-7
Gore	8	26	E	Passive copper	IBN4600-8
Gore	8	30	E	Passive copper	IBN4300-8
Gore	10	24	E	Passive copper	IBN4000-10
Leoni	6	26	E	Passive copper	L5590-A499-A59
Leoni	10	25	E	Passive copper	L5590-A499-A63
Meritec	2	28	N/E	Passive copper	9853041ABCX02.0
Meritec	5	28	N/E	Passive copper	9853041ABCX05.0
Meritec	10	24	N/E	Passive copper	944966
Gore	15	30	E	Active copper	IBN5320-15
Gore	20	28	E	Active copper	IBN5820-20
Gore	25	26	E	Active copper	IBN5620-25
Gore	30	24	E	Active copper	IBN5020-30

Active Optical Cables

An active optical cable consists of one optical cable assembly and two Optical Media Converters (OMCs).

The recommended OMC is the Emcore QTR3500 SmartLink™ DDR OMC.

In turn, an optical cable assembly consists of one parallel optical cable and two MPO (Multi-fiber Push On) male connectors, one at each end of the optical cable.

Because the optical port on the QTR3500 OMC is a male MPO connector, it is essential that any optical cable plugged into the OMC optical port be an MPO female connector plug as specified in the FOCIS 5 P-12-1-0-2-2 industry specification.

Index

A

access to RPM files, 27

B

back panel (IB NEM), 1

C

cable length, maximum, 2

changing the I/O configuration, 9

commands

ibnetdiscover, 33

ibstat, 33, 53

ifconfig, 47

lsmod, 46

lspci, 19

modprobe, 46

mstflint, 52

openibd, 32

openibd status, 32

ping, 48

route, 48

show /ch, 16

tar, 28

vstat, 42

Component Management page (ILOM), 14

Component Management Status page (ILOM), 15

configuring IPoIB, 30

controlling basic IB behavior, 32

D

DHCP, patch required for default IPoIB, 30

disk space required for installation (Windows), 35

F

features (IB NEM), 2

filler panel, installing to prevent overheating, 10

H

handling instructions, 9

hardware requirements for WinIB, 35

host interface, 2

hot-plug operations (supported), 2

hot-remove request, blade does not respond, 20

I

I/O configuration (changing), 9

IB cable length, maximum, 2

IB NEM

(installed), graphic view, 13

handling instructions, 9

how to release from blade, 21

LED status, 18

IB uplink interface, 2

ibnetdiscover command, 33

ibstat command

status details, 33

verifying firmware version on Linux, 53

ifcfg-ibn file

configure IBoIP, 29

creating/editing, 49

ifconfig command, 47

ifconfig command (check for network interfaces), 47

ILOM Component Management page, 14

ILOM Component Management Status page, 15

InfiniBand

- changing startup behavior, 48
- configuration file (openib.conf), 32
- controlling basic behavior, 32
- transfer rate, 2

InfiniBand drivers

- required for Windows platforms, 35

InfiniBand Trade Association (IBTA) interoperability

- Version 1.2 (supported), 2

interface (host), 2

interoperability, IBTA Version 1.2, 2

IPoIB configuration files, 29

K

KickStart configuration file (RHEL), 27

L

LED status (IB NEM), 18

Linux platforms (tested with OFED 1.0), 26

lsmod command, 46

lspci command (NEM visibility to OS), 19

M

Marketing part number, 1

maximum PCIe transfer rate, 2

Mellanox InfiniHost III Lx (MT25204) HCA, 2

memory (onboard), 2

modprobe command, 46

mstflint command (updating firmware on Linux), 52

N

NEM-to-blade connections

- ILOM CLI interface, 22

- ILOM web interface, 20

NEM-to-blade connections, determining, 20

network interface startup file (ifcg-ibn), 49

O

OFED

- configuring IPoIB interfaces, 30
- including IPoIB configuration files, 29
- installation directory (default), 30
- package components, 26
- required files, 28

OFED1.0.tgz, 28

onboard memory, 2

openib.conf file

- full path, 48

openib.conf file (InfiniBand configuration file), 32

openibd command, 32

operating system support, 3

overheating NEM modules, preventing, 10

P

PCI hot-remove request, blade does not respond, 20

PCIe transfer rate, 2

ping command (viewing the subnet), 48

platform support, 3

power, 2

product features, 2

R

Red Hat Package Manager files, access, 27

required Infiniband drivers for Windows, 35

RHEL4 U4 kernel support for HCA hardware, 23

RoHS, fully compliant, 2

route command (viewing subnet availability), 48

RPM, access to Red Hat Package Manager files, 27

S

software requirements for WinIB, 35

software stack, OFED, 26

startup behavior (IB), changing, 48

startup file, InfiniBand network interface, 49

state of the NEM-to-blade connections, determining, 20

T

tar command, 28

transfer rate (PCIe), 2

U

updating firmware on Linux with mstflint command, 52

upgrading IB drivers on Windows, 34

V

- verifying firmware version on Linux (ibstat command), 53
- vstat command (verifying node information), 42

W

- Windows platforms, 35
- WinIB
 - hardware requirements, 35
 - software requirements, 35

