



Sun™ Dual Port 4x DDR IB Host Channel Adapter PCIe® User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 820-3523-10
January 2008, Revision A

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

Copyright 2008 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 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.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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, Java, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ 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.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, États-Unis. Tous droits réservés.

Sun Microsystems, Inc. possède les droits de propriété intellectuelle relatifs à la technologie décrite dans ce document. En particulier, et sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs des brevets américains listés sur le site <http://www.sun.com/patents>, un ou les plusieurs brevets supplémentaires ainsi que les demandes de brevet en attente aux États-Unis et dans d'autres pays.

Ce document et le produit auquel il se rapporte sont protégés par un copyright et distribués sous licences, celles-ci en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Tout logiciel tiers, sa technologie relative aux polices de caractères, comprise, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, Java, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DÉCLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES DANS LA LIMITE DE LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.



Contents

Preface	xi
1. Overview	1
IB-HCA Card	2
I/O Interfaces	3
InfiniBand Interface	3
PCI Express Interface	3
LED Assignment	4
I ² C Compatible Interface	5
Power	5
Memory	6
System Memory	6
2. IB-HCA Card Installation	7
Installing the IB-HCA Card	7
▼ To Install the Adapter	7
▼ To Verify the Installation	9
3. InfiniBand Software on Linux	11
InfiniBand Software for Linux	11
OpenFabrics Enterprise Distribution for Linux	13

Installing the InfiniBand Drivers on Linux	14
▼ To Install IB Drivers From the Linux Distribution Source	15
▼ To Install the OFED Package	16
▼ To Verify Driver Installation on Linux	22
4. Internet Protocol Over InfiniBand	25
Configuring IPoIB on Linux	25
▼ To Install IPoIB Driver	26
▼ To Change IPoIB Configuration Without Rebooting	27
▼ To Change IB Startup Behavior on Linux	28
5. Updating IB EM Firmware	31
Updating IB EM Firmware for Linux	31
▼ To Update IB EM Firmware for Linux	32
A. Specifications for the Dual Port 4x DDR IB PCIe Host Channel Adapter	35
B. Instructions for Replacing a Tall Bracket with a Short Bracket on IB-HCA card	37
Removing a Tall Bracket and Assembling a Short Bracket	38
▼ To Assemble a Short Bracket	41
C. Diagnostic Software and Troubleshooting	45
SunVTS Diagnostic Testing	45
Using the SunVTS <code>ibhctest</code>	46
▼ To Use the <code>ibhctest</code> Command	46
Troubleshooting Tasks	47
Other Useful Utilities	47
<code>cfgadm</code>	48
<code>snoop</code>	48
<code>netstat</code>	49

kstat 49

Index 51

Figures

FIGURE 1-1	Dual Port 4x DDR IB PCIe Host Channel Adapter With Tall Bracket	2
FIGURE 1-2	I/O Panel With Dual Ports and LEDs (Tall Bracket)	5
FIGURE 2-1	Adapter Connected to Two Sun IB Switches	9
FIGURE B-1	Tall Bracket Of The Dual Port 4x DDR IB PCIe Host Channel Adapter	37
FIGURE B-2	Connector Retention Clip	38
FIGURE B-3	Extracting Connector Clip	39
FIGURE B-4	Unscrew Bracket Screws	39
FIGURE B-5	Removing the Tall Bracket	40
FIGURE B-6	Place Short Bracket onto Card	41
FIGURE B-7	Attach Bracket onto Card using Screws	42
FIGURE B-8	Sliding Connector Clip Evenly	42
FIGURE B-9	Fix Clip Hooks into Place Using Screwdriver	43
FIGURE B-10	Assembled Short Bracket View	43

Tables

TABLE 1-1	LEDs	4
TABLE 3-1	RHEL AS 4-U4 Packages	12
TABLE 3-2	Required Packages for InfiniBand Support	13
TABLE 3-3	Linux Platforms and OFED Release	14
TABLE A-1	Specifications For The Dual Port 4x DDR IB PCIe Host Channel Adapter (SDR)	35
TABLE A-2	Specifications For The Dual Port 4x DDR IB PCIe Host Channel Adapter (DDR)	36

Preface

The *Sun Dual Port 4x DDR IB Host Channel Adapter PCIe User's Guide* provides an overview, installation instructions and adapter specifications of the Sun™ Dual Port 4x DDR IB PCIe Host Channel Adapter. The instructions in this guide are designed for system administrators with experience installing network hardware and software.

Note – The Sun Dual Port 4x DDR IB PCIe Host Channel Adapter is based on Mellanox Technologies' MT25408 ConnectX™ IB adapter (HCA) device.

How This Document Is Organized

[Chapter 1](#) provides an overview of the Dual Port 4x DDR IB PCIe Host Channel Adapter.

[Chapter 2](#) provides installation instructions and safety warnings for the adapter.

[Chapter 3](#) provides InfiniBand Software overview

[Chapter 4](#) describes configuring Internet Protocol over InfiniBand.

[Chapter 5](#) provides information on updating IB ExpressModule firmware.

[Appendix A](#) provides specifications for the Low Profile Adapter models.

[Appendix B](#) provides instructions for replacing brackets for the Low Profile Adapter.

[Appendix C](#) provides instruction on diagnostics and troubleshooting.

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
- 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> .

Note – Characters display differently depending on browser settings. If characters do not display correctly, change the character encoding in your browser to Unicode UTF-8.

Documentation, Support, and Training

Sun Function	URL
Documentation	http://docs.sun.com
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.

PCI Express specifications can be found at this web site:

<http://www.pcisig.com>

Mellanox Technologies HCA information can be found at this web site:

<http://www.mellanox.com>

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/documentation/feedback>

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

Sun Dual Port 4x DDR IB Host Channel Adapter PCIe User's Guide, part number 820-3523-10.

Overview

This document provides detailed procedures for installing the hardware and software for the Dual Port 4x DDR IB PCIe Host Channel Adapter card. The cards described in this guide have the following main features:

- IBTA (InfiniBand Trade Association) v1.2 compliant
- Two 4x InfiniBand copper ports for connecting InfiniBand traffic (4x IB connectors)
- PCI Express expansion board with an x8 edge connector compatible to the PCI Express 1.0a specification
- 'Media detect circuit' supporting external InfiniBand fiber solutions
- European Union Restriction of Hazardous Substances (RoHS) compliant


The boards differ in

- 4x IB port speed support: 10 Gb/s (SDR-Single Data Rate) or 20 Gb/s (DDR-Double Data Rate)
- Bracket height: short or tall

Note – The PCI Express (or PCIe) interface is sometimes called PCI-E.

IB-HCA Card

FIGURE 1-1 Dual Port 4x DDR IB PCIe Host Channel Adapter With Tall Bracket



I/O Interfaces

Each IB-HCA card includes the following interfaces:

- Two 4X InfiniBand copper connectors
- PCI Express x8 edge connector
- I/O panel LEDs

InfiniBand Interface

The Dual Port 4x DDR IB PCIe Host Channel Adapter device is compliant with the *InfiniBand Architecture Specification, Release 1.2*. It has two compliant 4X InfiniBand ports, A and B. The IB-HCA card provides access to these ports by means of two 4X InfiniBand connectors for external InfiniBand copper cables also compliant with the *InfiniBand Architecture Specification, Release 1.2*. Connector 1 connects to port A of the device, while connector 2 connects to port B.

Furthermore, the IB-HCA card is embedded with a ‘media detect circuit’ that supports external InfiniBand fiber solutions. These are connected to the InfiniBand port connectors using active media converters, such as Emcore’s *QTR3400 Smart Module* or Fujitsu’s *FPD-010R008-0E o-microGiGaCN*.

PCI Express Interface

The Dual Port 4x DDR IB PCIe Host Channel Adapter device has eight Tx/Rx pairs of SerDes providing for a PCI Express x8 interface, version 1.0a compatible. The device can be either a master initiating the PCI Express bus operations or a slave responding to PCI bus operations. The PCI Express bus can connect to either a host CPU in an IB-HCA application or to an I/O device (such as Gigabit Ethernet) when used as a Target Channel Adapter.

LED Assignment

The board has four LEDs located on the I/O panel - 2 LEDs per 4X port. See [TABLE 1-1](#).

TABLE 1-1 LEDs

Port Number	LED Name
Port 1	Physical Link - Green
	Data Activity - Amber
Port 2	Physical Link - Green
	Data Activity - Amber

The physical link (green) illuminates once VAPI (InfiniBand Verbs API) is started and a physical connection is made between two nodes. The data activity link (amber) illuminates once the InfiniBand network is discovered over the physical link. The activity link is a steady yellow when it is discovered but no data is being passed. The activity link blinks when data is being passed. If the LEDs are not active, either the physical or the logical (or both) connections have not been established. See [FIGURE 1-2](#).

Short brackets have the same port and LED footprints as tall brackets. See [Appendix B](#).

FIGURE 1-2 I/O Panel With Dual Ports and LEDs (Tall Bracket)




Figure Legend Ports and LEDs of the ExpressModule

-
- 1 InfiniBand Port 1
 - 2 Physical Link (Green LED)
 - 3 Data Activity Link (Amber LED)
 - 4 InfiniBand Port 2
 - 5 Data Activity Link (Amber LED)
 - 6 Physical Link (Green LED)
-

I²C Compatible Interface

A three-pin header, designated with reference name J5 on the IB-HCA card, is provided as the I²C compatible interface.

Power

All adapter cards receive power from the PCI Express Edge connector. All other required power voltages are generated by on-board switch mode regulators. For power consumption see [Appendix A](#).

Memory

The IB-HCA card support multiple memory devices through the PCI Express, Flash, and I²C-compatible interfaces.

System Memory

The IB-HCA card utilizes the PCI Express interface to store and access IB fabric connection information on the system memory.

IB-HCA Card Installation

The IB-HCA card is a standard PCI Express x8 card with a standard x8 edge connector. Consult the host machine documentation for instructions on how to install a PCI Express card

Installing the IB-HCA Card



Caution – Electronic components on printed circuit boards are extremely sensitive to static electricity. Ordinary amounts of static electricity generated by your clothing or work environment can damage the electronic equipment. When installing the IB-HCA in a system, use antistatic grounding straps and antistatic mats to help prevent damage due to electrostatic discharge.

▼ To Install the Adapter

Refer to your system installation or service manual for detailed instructions for the following steps:

1. **Power off your server, using the standard shutdown procedures described in your system service manual.**

Documentation is available on the <http://docs.sun.com> web site.

2. **Remove the cover from the unit to access the card slots and connectors.**
3. **Select an available PCIe x8 slot and remove the blank front panel.**

If you are replacing an existing card, remove the card.

4. **Install the IB-HCA card into the slot, pushing the card's edge connector into the connector on the chassis.**

Ensure that the front plate on the IB-HCA card mounts flush with the chassis panel opening.

5. **If applicable, install the screw in front plate to secure the IB-HCA card into the chassis.**
6. **Connect the 4x end of the InfiniBand I/O cable(s) to the IB-HCA port connector(s).**

Ensure that the connectors are properly engaged.



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


7. **Replace the cover on the unit.**
8. **If not already connected, connect the 12x end of the InfiniBand I/O cables to the appropriate ports on the switch or switches.**

The IB-HCA ports can be connected to different ports on the same switch or to a port on different switches. See [FIGURE 2-1](#) for an example of cabling an adapter to two switches.

9. **Turn power back on to the system and allow the server to reboot.**

This completes the hardware installation. Proceed to the verification instructions in [“To Verify the Installation” on page 9](#).

FIGURE 2-1 Adapter Connected to Two Sun IB Switches



▼ To Verify the Installation

The InfiniBand switch should automatically recognize the IB-HCA when it is connected to the fabric if the IB Subnet Manager is running on the switch.

1. **Ensure that the cables are connected to the adapter and switches.**
2. **Verify that the IB Subnet Manager is running on the IB switch.**
Refer to the *Sun IB Switch 9P Administration Guide* (819-0502) manual for information on the IB Subnet Manager.
3. **Check that the green LED is illuminated for each port that is connected to the switch.**

If the green LED is not on, check the cable connections at the adapter and at the switch.

4. Check that the amber LED is illuminated for each port that is connected to the switch.

If the amber LED is not on, refer to [Appendix C, “Diagnostic Software and Troubleshooting”](#) on page 45 for more information.

5. Verify that the IB-HCA ports are up and the driver is attached by typing:

```
# dmesg | grep tavor
```

The output shows system diagnostic messages that have the string `tavor` in the message. Included in the output is a message that tells whether the port is up or down. For example:

```
Feb 17 20:30:42 hostname px_pci: [ID 370704 kern.info] PCI-device:
pciex15b3,6278@0, tavor0
Feb 17 20:30:42 hostname genunix: [ID 936769 kern.info] tavor0 is
/pci@7c0/pci@0/pci@8/pciex15b3,6278@0
Feb 17 20:30:42 hostname tavor: [ID 548581 kern.info] tavor0: FW ver:
0004.0007.0190, HW rev: a0
Feb 17 20:30:42 hostname tavor: [ID 827127 kern.info] tavor0: MT23108
Sun Microsystems InfiniBand HCA (0x0003ba0001001c18)
Feb 17 20:30:45 hostname genunix: [ID 979957 kern.info] NOTICE: tavor0:
port 1 up
Feb 17 20:30:45 hostname genunix: [ID 611667 kern.info] NOTICE: tavor0:
port 2 up
Feb 17 20:30:58 hostname genunix: [ID 408114 kern.info]
/pci@7c0/pci@0/pci@8/pciex15b3,6278@0 (tavor0) online
```

InfiniBand Software on Linux

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 (the ExpressModule) and an InfiniBand software stack.

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

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.

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 and OFED Release 1.2.5 or later, you do not need to install or configure additional drivers to support the IB ExpressModule (IB EM).

Specifically, RHEL AS 4-U4 contains support in the kernel for IB-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. The InfiniBand modules 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 described in [TABLE 3-1](#).

TABLE 3-1 RHEL AS 4-U4 Packages

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

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

The packages selected to support any given configuration will vary. [TABLE 3-2](#) lists the packages considered the absolute minimum needed to support the environment described in this guide.

TABLE 3-2 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
OFED Release 1.2.5 or later	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. You can also download the required files from OpenFabrics.org. For information on both of these procedures, see [“Installing the InfiniBand Drivers on Linux”](#) on page 14.

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

TABLE 3-3 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 1.2.5 of the OpenFabrics stack. Note: RHEL AS 4-U4 includes an older version of OFED, so you must install OFED Release 1.2.5 or a later version.
SLES9 SP3 or later, SLES10	Sun has tested OFED Release 1.2.5 for the SLES10 platform. Note: You must have OFED Release 1.2.5 or a later version.

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:

- [“To Install IB Drivers From the Linux Distribution Source”](#) on page 15
- [“To Install the OFED Package”](#) on page 16

If you need to determine whether or not the drivers are already installed, see [“To Verify Driver Installation on Linux”](#) on page 22.

▼ To Install IB Drivers From the Linux Distribution Source

1. Obtain the Red Hat Package Manager (RPM) files containing the InfiniBand drivers.

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.

2. 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
.
.
.
```

3. If you are running the CSH or TCSH shell, enter the `rehash` command to rebuild the shell's view of available executables.
4. Enter the `ibstat` command to verify that the OS sees the IB EM.

```
> 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
```

5. (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 [“To Verify Driver Installation on Linux” on page 22](#).
6. (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 [“To Install IPoIB Driver” on page 26](#).

▼ To Install the OFED Package

Note – The Dual Port 4x DDR IB PCIe Host Channel Adapter requires OFED Release 1.2.5 or later.

1. On the Sun Blade Server Module, log in as `root` and copy the required files from the following location:

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

In the following example, `OFED-1.2.5.tar` is used only as an example.

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

2. From root, extract the files by typing:

```
> tar -zxvf OFED-1.2.5.tar
```

3. From the OFED-1.2.5 directory, initiate the installation process by typing:

```
> ./install.sh
```

4. When the InfiniBand OFED Distribution Software Installation Menu appears, select option 2 (Install OFED Software).
5. When the Select OFED Software menu appears, select 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]: n
```

7. Next, you are asked if you wish to create/install an openmpi RPM with gcc. Again, type 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]: n
```

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]: Y
```

8. **Type Y to continue, as shown above.**

Next, you are prompted to configure InfiniBand IP support.

9. **Type Y when asked if you want to include IPoIB configuration files.**

```
Do you want to include IPoIB configuration files (ifcfg-ib*)?  
[Y/n]: Y
```

10. **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]:
```

11. 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 ...
```

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.

Assuming the IB EM hardware is installed (and, therefore, an InfiniBand HCA is present), you are prompted to configure InfiniBand IP support.

12. Enter Y in response to the following prompt:

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

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.

The system next displays the current configuration.

13. When asked if you want to change the configuration as displayed, type *y*.

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

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]: Y
```

14. Type *Y* to save the configuration.

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

15. Iterate the InfiniBand configuration over all InfiniBand interfaces.

Enter a valid IP configuration for each network interface.

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

16. Enter n to complete this part of the installation.

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

You should see a message like the following.

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

17. Press Enter.

The InfiniBand OFED Distribution Software Installation Menu is displayed.

18. Type 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 option
```

where *option* can be `start`, `stop`, or `status`.

19. After a successful installation, reboot the Server Module.

After the reboot, the Server Module should come up as a functional member of the InfiniBand fabric.

▼ To Verify Driver Installation on Linux

1. **Verify that the Linux software driver is installed and attached to the IB EM by typing the `openibd status` command.**

When using the `openibd` command, type the entire path as shown in this example..

```
> /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
.
.
.
```

This 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`. To view details of operational status, type 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
sysimgguid=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
sysimgguid=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
```


Internet Protocol Over InfiniBand

This chapter describes configuration aspects of running the Internet Protocol over InfiniBand (IPoIB) on Linux.

Configuring IPoIB on Linux

You might decide to change your IPoIB configuration for a variety of reasons, including the installation of an additional InfiniBand ExpressModule (IB EM).

▼ To Install IPoIB Driver

1. **Determine whether the IPoIB driver is already installed by typing 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
findex         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. Type 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 [“To Change IPoIB Configuration Without Rebooting” on page 27](#).

▼ To Change IPoIB Configuration Without Rebooting



Caution – The changes made by this procedure are only temporary. These changes are lost during a reboot. To make permanent changes, see [“To Change IB Startup Behavior on Linux” on page 28](#).

1. Start the InfiniBand IP network by typing 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` is unconfigured, it appears 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. Type 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  Ifac
e
10.0.0.0   *                255.255.255.0    U      0       0    0    ib0
10.8.134.0 *                255.255.255.0    U      0       0    0    eth0
169.254.0.0 *               255.255.0.0     U      0       0    0    eth0
default   ban3rtr0d0      0.0.0.0          UG     0       0    0    eth0
```

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.

▼ To Change IB Startup Behavior on Linux

You can change the InfiniBand startup behavior by editing the `openib.conf` configuration file. (The full path is `/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`.

The following is an example of the `ipoib.conf` file.

```
# 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 startup file (`ifcfg-ibn`) in your startup scripts directory.

The directory for startup scripts might be `/etc/sysconfig/network/`, `/etc/sysconfig/network-scripts/`, or similar, depending on your version of Linux distribution.

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 automatically create the `ifcfg-ibn` configuration files for all IB network interfaces present. If you install an IB EM *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 EM Firmware

This chapter provides information on updating the IB ExpressModule (IB EM) firmware on Linux.

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

Updating IB EM Firmware for Linux

The firmware version on your IB EM should be ready to use and should not require updating. However, if you want to update the firmware, you must 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 EM 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 EM in the PCI `bus:dev.fun` format to satisfy the `mstflint` command `-d device` syntax requirement.

▼ To Update IB EM Firmware for Linux

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

```
> 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)
```

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

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

```
> mstflint -d 81:00.0 -i ibem.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 that 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, type `y`.

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

4. After resetting the Server Module (or the EM), 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
```


Specifications for the Dual Port 4x DDR IB PCIe Host Channel Adapter

TABLE A-1 lists the specifications for the IB-HCA card with IB running at SDR.

TABLE A-1 Specifications For The Dual Port 4x DDR IB PCIe Host Channel Adapter (SDR)

Physical	
Size:	2.2in. x 5.4in. (5.5cm x 13.7cm)
Air Flow:	200LFM @55°C
4X 10Gb/s Connector:	InfiniBand (Copper, current rating: 0.5A max) with active media adapter support
Protocol Support	
InfiniBand:	IBTA v1.2, Auto-Negotiation 10Gb/s, 2.5Gb/s
QoS:	8 InfiniBand Virtual Lanes for each port
RDMA Support:	Yes, All Ports
Power and Environmental	
Voltage:	12V, 3.3V
Maximum Power:	11.5W
Temperature:	0°C to 55°C
Regulatory	
Safety:	IEC/EN 60950-1:2001, ETSI EN 300 019-2-2
Environmental:	IEC 60068-2- 64, 29, 32
RoHS:	RoHS-R5

TABLE A-2 lists the specifications for the IB-HCA card with IB running at DDR.

TABLE A-2 Specifications For The Dual Port 4x DDR IB PCIe Host Channel Adapter (DDR)


Physical	
Size:	2.2in. x 5.4in. (5.5cm x 13.7cm)
Air Flow:	200LFM @55°C
4X 20Gb/s Connector:	InfiniBand (Copper, current rating: 0.5A max) with active media adapter support
Protocol Support	
InfiniBand:	IBTA v1.2, Auto-Negotiation* (20Gb/s, 5Gb/s) or (10Gb/s, 2.5Gb/s)
QoS:	8 InfiniBand Virtual Lanes for each port
RDMA Support:	Yes, All Ports
Power and Environmental	
Voltage:	12V, 3.3V
Maximum Power:	12W
Temperature:	0°C to 55°C
Regulatory	
EMC:	FCC 47 CFR part 15:2005, subpart B, class A ICES-003:2004 Issue 4, class A VCCI V-3/2005.04, class A EN 55022:1998+A1:2000+A2:2003 class A, EN 61000-3-2:2000+A2:2005, EN 61000-3-3:1995+A1:2001, EN 55024:1998 + A1:2001+A2:2003 standards, harmonized under EMC Directive 89/336/EEC
Safety:	IEC/EN 60950-1:2001 ETSI EN 300 019-2-2
Environmental:	IEC 60068-2- 64, 29, 32
RoHS:	RoHS-R5

* The auto-negotiation protocol is proprietary of Mellanox Technologies and compliant with the *InfiniBand Architecture Specification, Release 1.2*.

Instructions for Replacing a Tall Bracket with a Short Bracket on IB-HCA card

FIGURE B-1 shows the bracket-side view of a dual-port IB-HCA card.

FIGURE B-1 Tall Bracket Of The Dual Port 4x DDR IB PCIe Host Channel Adapter



Removing a Tall Bracket and Assembling a Short Bracket

The connector retention clips are used to hold the short or tall bracket in place on the card. Before the bracket can be replaced, the connector retentions clips have to be removed. [FIGURE B-2](#) shows a connector retention clip and the designated names of its sections.

FIGURE B-2 Connector Retention Clip




Figure Legend Parts of the Retention clip

-
- | | |
|---|-------------|
| 1 | Clip top |
| 2 | Clip center |
| 3 | Clip hook |
-

1. Remove Connector Clips

- Using a small flat head screwdriver, gently push up one hook of a connector's clip toward the connector's top side as shown in [FIGURE B-3 \(1\)](#).
- Push the other hook each of the two clip's hook towards the connector's top side - as shown in [FIGURE B-3 \(2\)](#).
- Finally, pull the clip away from its center - as shown in [FIGURE B-3 \(3\)](#).

FIGURE B-3 Extracting Connector Clip




d. Repeat the above actions for the second connector's clip.

2. Unscrew Bracket Screws

a. Unscrew both screws from the card using a torque screwdriver as shown in [FIGURE B-4](#).


FIGURE B-4 Unscrew Bracket Screws




3. Detach the Bracket

- a. Grip the bracket as shown in [FIGURE B-5](#) placing your thumb on the LED component. In a rotating move toward the component side of the card, slide the bracket out of the connectors.
- b. Gently hold your thumb on the LED component as shown in [FIGURE B-5](#), bubble 2. At the same time extract the bracket, while making sure to protect the LEDs.
- c. Detach the tall bracket in a rotating move as shown in [FIGURE B-5](#), bubble 3.


FIGURE B-5 Removing the Tall Bracket



Card with Clips and Screws Removed



Grip the Card and Bracket




Rotate the Bracket Toward the Component Side to Remove

▼ To Assemble a Short Bracket

1. Place Short Bracket onto Card

- a. Gently place the bracket onto the card fitting the connectors through the bracket connector holes. Make sure the LEDs are aligned into their intended bracket holes.

FIGURE B-6 Place Short Bracket onto Card



2. Attach Short Bracket to Card

- a. Insert a screw along with a washer into each of the two holes on the card intended for holding the bracket as shown in [FIGURE B-7](#). Use a torque screwdriver to apply up to 2 lbs-in torque on each screw.


FIGURE B-7 Attach Bracket onto Card using Screws



3. Install Connector Clips


- a. Gently push one clip onto the connector. Make sure to slide both clip hooks (sides) around the connector evenly as shown in [FIGURE B-8](#).

FIGURE B-8 Sliding Connector Clip Evenly




- b. Use a small flat head screwdriver to gently slide the clip's hook towards the connector's base side as shown in [FIGURE B-9](#).

FIGURE B-9 Fix Clip Hooks into Place Using Screwdriver



- c. Repeat this step for the second clip. See [FIGURE B-10](#).

FIGURE B-10 Assembled Short Bracket View



Diagnostic Software and Troubleshooting

This appendix provides an overview of the SunVTS™ diagnostic application and troubleshooting tips for checking the adapter. This is for the Solaris OS only and not for Linux. There is also a section outlining some common troubleshooting issues. This appendix contains the following sections:

- “SunVTS Diagnostic Testing” on page 45
- “Troubleshooting Tasks” on page 47
- “Other Useful Utilities” on page 47

SunVTS Diagnostic Testing

The SunVTS software executes multiple diagnostic hardware tests from a single user interface, and is used to verify the configuration and functionality of most hardware controllers and devices. The SunVTS software operates primarily from a graphical user interface, enabling test parameters to be set quickly and easily while a diagnostic test operation is being performed.

Note – SunVTS diagnostic software is currently available for x86/x64 systems in the Solaris™ 10 Operating Systems.

Refer to the SunVTS section in your Solaris 10 documentation for instructions on how to run and monitor the `nettest` diagnostic. These SunVTS documents are available online at the following URL:

<http://docs.sun.com/app/docs/prod/test.validate/sunvts/index.html>

Select the document for the SunVTS release on your system

Using the SunVTS `ibhctest`

The `ibhctest` diagnostic test checks the functionality of Dual Port 4x DDR IB PCIe Host Channel Adapter card. This test can be run from the SunVTS user interface, or it can be run from the command line. See the *SunVTS 6.0 Test Reference Manual* (817-7665) for more information about the `ibhctest` test.

The `ibhctest` diagnostic test is included in the SunVTS 6.0 Patch Set 1 and subsequent SunVTS software releases. SunVTS 6.0 Patch Set 1 is available for downloading from the SunSolveSM web site <http://sunsolve.sun.com> using the following patch numbers:

- SPARC: patch 118962-01, or higher
- X86: patch 118961-01, or higher

The adapter and Tavor device driver must be installed, and the IB port interface must be configured `offline` for the `ibhctest` to run. A loopback cable is not needed because `ibhctest` includes an internal loopback test. Use the following procedure when running the `ibhctest` command. See more details in your in the documentation for the version of SunVTS you are using.

▼ To Use the `ibhctest` Command

1. **Ensure that the SunVTS software and the Tavor driver are installed on your system, by typing:**

```
# pkginfo SUNWvts SUNWvtsx SUNWtavor
```

If a SunVTS software package is not installed, refer to the *SunVTS User's Guide* for installation instruction. If the `SUNWtavor` package is not installed, check your Solaris Operating System documentation for software package information.

2. **Unplumb the interface from the system, using the `ifconfig` command:**

```
# ifconfig ibdn down unplumb
```

where *n* is the instance number of the interface.

3. **Refer to your SunVTS documentation for instructions on how to run the `ibhctest` command.**

Troubleshooting Tasks

The following tasks can be useful when troubleshooting the IB-HCA and the link.

- **Verify that the InfiniBand software packages are installed on the host.** To do this, type the command:

```
pkginfo -l | egrep "InfiniBand | Tavor | udapl"
```

Check that the following packages are installed:

- SUNWib – Sun InfiniBand Framework
- SUNWtavor – Sun Tavor HCA Driver
- SUNWipoib – Sun IP over InfiniBand
- SUNWudapl_r – Direct Access Transport (DAT) registry package (*root*)
- SUNWudapl_u – Direct Access Transport (DAT) registry packages (*usr*)
- SUNWudapl_{tr} – Service Provider for Tavor packages (*root*)
- SUNWudapl_{tu} – Service Provider for Tavor packages

If an InfiniBand software package is not installed, check your Solaris Operating System documentation for software package information.

- **Check the system log or console for Tavor driver error messages.**

See `tavor(7D)` for error messages and descriptions. When the driver is attached to a port on the adapter, the following message is sent.

```
tavorn: port m up (link width 4x).
```

In the message, *n* is the instance of the Tavor device number and *m* is the port number on the adapter.

One way to check Tavor messages is by typing the following command:

```
# dmesg | grep tavor
```

Other Useful Utilities

These utilities can display status and other information about InfiniBand devices:

- `cfgadm`

- snoop
- netstat
- kstat

cfgadm

The `cfgadm` utility displays status and other information about the IB-HCA and IB fabric. See `cfgadm_ib(1M)` for details. For example:

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
hca:21346543210a987	IB-HCA	connected	configured	ok
ib	IB-FABRIC	connected	configured	ok
ib::80020123456789a	IB-IOC	connected	configured	ok
ib::802abc9876543	IB-IOC	connected	unconfigured	unknown
ib::80245678,ffff,ipib	IB-VPPA	connected	configured	ok
ib::12245678,0,nfs	IB-PORT	connected	configured	ok
ib::21346543,0,hnfs	IB-HCA_SVC	connected	configured	ok
ib::sdp,0	IB-PSEUDO	connected	configured	ok

snoop

The `snoop` program captures and inspects network packets. See the `snoop(1M)` man page for details. For example:

```
# snoop -d ibd1
```

```
Using device /dev/ibd1 (promiscuous mode)
```

ib-1-167 -> *	ARP C Who is 199.1.1.168, ib-1-168 ?
ib-1-168 -> ib-1-167	ARP R 199.1.1.168, ib-1-168 is
0:2:4:7:0:0:0:0:a:4:7c:4f:0:2:c9:2:0:0:55:91	
ib-1-167 -> ib-1-168	ICMP Echo request (ID: 35608 Sequence number: 0)
ib-1-168 -> ib-1-167	ICMP Echo reply (ID: 35608 Sequence number: 0)

netstat

netstat shows network status. See the netstat(1M) man page for details. For example:

```
# netstat -I ibd 4
```

input		ibd1			output			input (Total)		output		
packets	errs	packets	errs	colls	packets	errs	packets	errs	packets	errs	colls	
2458394	0	2458268	0	0	2467288	0	2465951	0	0			
92233	0	92237	0	0	92247	0	92238	0	0			
92703	0	92702	0	0	92709	0	92704	0	0			

kstat

kstat displays kernel statistics. See the kstat(1M) man page for details. For example:

```
# kstat ibd:1
```

module:	ibd	instance:	1
name:	ibd1	class:	net
			0
opackets			27381595
opackets64			27381595
promisc		off	xmt_badinterp
	0		
xmtretry			4

Index

A

- access to RPM files, 15
- adapter
 - cabling, 8
 - power, 5
 - specifications, 35

B

- Bracket
 - height, 1
 - replacing, 37

C

- commands
 - ibnetdiscover, 23
 - ibstat, 22, 33
 - ifconfig, 27
 - lsmod, 26
 - modprobe, 26
 - mstflint, 32
 - openibd, 21
 - openibd status, 22
 - ping, 28
 - route, 28
- configuring IPoIB, 19
- controlling basic IB behavior, 21

D

- DHCP, patch required for default IPoIB, 19
- diagnostics, ibhcatetest, 46
- documentation
 - SunVTS, 45

I

- I/O
 - cables, 8
 - interfaces, 3
 - panel, 5
- IB Subnet Manager, 9
- IB-HCA
 - installation, 8
- ibhcatetest diagnostics, 46
- ibnetdiscover command, 23
- ibstat command
 - status details, 22
 - verifying firmware version on Linux, 33
- ifcfg-ibn file
 - configure IBoIP, 18
 - creating/editing, 30
- ifconfig command, 27
- ifconfig command (check for network interfaces), 27
- InfiniBand
 - changing startup behavior, 28
 - configuration file (openib.conf), 21
 - controlling basic behavior, 21
- IPoIB configuration files, 18

K

- KickStart configuration file (RHEL), 15

L

- LED
 - assignment, 4
 - location, 5
 - verification, 9

LEDs
 verification, 9
Linux platforms (tested with OFED 1.0), 14
lsmod command, 26

M
modprobe command, 26
mstflint command (updating firmware on Linux), 32

N
network interface startup file (ifcg-ibn), 30

O
OFED
 configuring IPoIB interfaces, 19
 including IPoIB configuration files, 18
 installation directory (default), 19
 package components, 14
openib.conf file
 full path, 28
openib.conf file (InfiniBand configuration file), 21
openibd command, 21, 22

P
Panel
 illustration, 5
panel
 LED, 3
ping command (viewing the subnet), 28

R
Red Hat Package Manager files, access, 15
RHEL4 U4 kernel support for HCA hardware, 11
route command (viewing subnet availability), 28
RPM, access to Red Hat Package Manager files, 15

S
software packages, 47
software stack, OFED, 14
startup behavior (IB), changing, 28
startup file, InfiniBand network interface, 30
Subnet Manager, 9
SunVTS
 diagnostic application, 45

documentation, 45
ibhctest diagnostics, 46

T
troubleshooting
 IB-HCA, 47
 tasks, 47
 utilities, 47

U
updating firmware on Linux with mstflint command, 32
utilities, 47
 cfgadm, 48
 kstat, 49
 netstat, 49
 snoop, 48

V
verifying firmware version on Linux (ibstat command), 33