



Sun™ Dual Port 4x IB Host Channel Adapter PCI-X Installation and User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 819-1280-10
September 2005, Revision A

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

Copyright 2005 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, SunVTS, Sun Fire, SunSolve, 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 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui 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.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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, Java, AnswerBook2, docs.sun.com, SunVTS, Sun Fire, SunSolve, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-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 Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation 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 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 licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" 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 CONTREFAÇON.

Contents

Regulatory Compliance Statements	xi
Declaration of Conformity	xv
Preface	xvii
How This Book Is Organized	xvii
Using UNIX Commands	xviii
Shell Prompts	xviii
Typographic Conventions	xix
Related Documentation	xix
Documentation, Support, and Training	xx
Third-Party Web Sites	xx
Contacting Sun Technical Support	xx
Sun Welcomes Your Comments	xxi
1. Product Overview	1
Product Features	1
Hardware Description	1
I/O Interfaces	3
InfiniBand Interface	3
PCI-X Interface	3

LED Displays	3
Hardware, Software, and Cable Requirements	5
2. Adapter Installation	7
IB-HCA Bus and Slot Selection	7
▼ Installing the Adapter	7
▼ Verifying the Installation	9
3. Configuring IP Over InfiniBand	11
IPoIB Configuration	11
▼ Configuring IPoIB	11
Managing IB Devices for IPoIB	13
4. InfiniBand Software Overview	15
InfiniBand Software	15
User Direct Access Programming Library	16
IP over InfiniBand Driver	16
InfiniBand Transport Framework	17
InfiniBand Management Framework	17
Tavor HCA Driver	17
Related Software Documentation	18
Manuals	18
Man Pages	18
A. InfiniBand Cable Pin Assignments	19
B. Diagnostic Software and Troubleshooting	23
SunVTS Diagnostic Testing	23
Using the SunVTS <code>ibhctest</code>	24
▼ Using the <code>ibhctest</code>	25
Troubleshooting Tasks	25

Other Useful Utilities 26

 cfgadm 26

 snoop 27

 netstat 27

 kstat 27

Index 29

Figures

FIGURE 1-1	Sun Dual Port 4x IB Host Channel Adapter PCI-X Card	2
FIGURE 1-2	LED Layout	4
FIGURE 2-1	Typical Adapter-to-Switch Configuration	9
FIGURE 4-1	InfiniBand Software Stack	16
FIGURE A-1	Sun 4x-to-12x InfiniBand Cable Pin Assignments	21

Tables

TABLE 1-1	IB-HCA LEDs	3
TABLE 1-2	Hardware, Software, and Cable Requirements	5
TABLE B-1	SunVTS Documentation	24

Regulatory Compliance Statements

Your Sun product is marked to indicate its compliance class:

- Federal Communications Commission (FCC) — USA
- Industry Canada Equipment Standard for Digital Equipment (ICES-003) — Canada
- Voluntary Control Council for Interference (VCCI) — Japan
- Bureau of Standards Metrology and Inspection (BSMI) — Taiwan

Please read the appropriate section that corresponds to the marking on your Sun product before attempting to install the product.

FCC Class A Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if it is not installed and used in accordance with the instruction manual, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

FCC Class B Notice

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/television technician for help.

Modifications: Any modifications made to this device that are not approved by Sun Microsystems, Inc. may void the authority granted to the user by the FCC to operate this equipment.

ICES-003 Class A Notice - Avis NMB-003, Classe A

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

ICES-003 Class B Notice - Avis NMB-003, Classe B

This Class B digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.


VCCI 基準について

クラス A VCCI 基準について

クラス A VCCI の表示があるワークステーションおよびオプション製品は、クラス A 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

クラス B VCCI 基準について

クラス B VCCI の表示  があるワークステーションおよびオプション製品は、クラス B 情報技術装置です。これらの製品には、下記の項目が該当します。

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス B 情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。取扱説明書に従って正しい取り扱いをしてください。

BSMI Class A Notice

The following statement is applicable to products shipped to Taiwan and marked as Class A on the product compliance label.

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

CCC Class A Notice

The following statement is applicable to products shipped to China and marked with "Class A" on the product's compliance label.

以下声明适用于运往中国且其认证标志上注有 "Class A" 字样的产品。

声明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其干扰采取切实可行的措施。



GOST-R Certification Mark



Declaration of Conformity

Compliance Model Number: **IB-HCA**
Product Family Name: **Sun Dual Port 4x IB Host Channel Adapter PCI-X**

EMC

USA—FCC Class A

This equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This equipment may not cause harmful interference.
2. This equipment must accept any interference that may cause undesired operation.

European Union

This equipment complies with the following requirements of the EMC Directive 89/336/EEC:

As Telecommunication Network Equipment (TNE) in Both Telecom Centers and Other Than Telecom Centers per (as applicable):

EN300-386 V.1.3.1 (09-2001) Required Limits:

EN55022/CISPR22	Class A
EN61000-3-2	Pass
EN61000-3-3	Pass
EN61000-4-2	6 kV (Direct), 8 kV (Air)
EN61000-4-3	3 V/m 80-1000MHz, 10 V/m 800-960 MHz, and 1400-2000 MHz
EN61000-4-4	1 kV AC and DC Power Lines, 0.5 kV Signal Lines
EN61000-4-5	2 kV AC Line-Gnd, 1 kV AC Line-Line and Outdoor Signal Lines, 0.5 kV Indoor signal Lines > 10m
EN61000-4-6	3 V
EN61000-4-11	Pass

As Information Technology Equipment (ITE) Class A per (as applicable):

EN55022:1998/CISPR22:1997 Class A

EN55024:1998 Required Limits:

EN61000-4-2	4 kV (Direct), 8 kV (Air)
EN61000-4-3	3 V/m
EN61000-4-4	1 kV AC Power Lines, 0.5 kV Signal and DC Power Lines
EN61000-4-5	1 kV AC Line-Line and Outdoor Signal Lines, 2 kV AC Line-Gnd, 0.5 kV DC Power Lines
EN61000-4-6	3 V
EN61000-4-8	1 A/m
EN61000-4-11	Pass
EN61000-3-2	Pass
EN61000-3-3	Pass

Safety: *This equipment complies with the following requirements of the Low Voltage Directive 73/23/EEC:*

EC Type Examination Certificates:

EN60950-1:2001, 1st Edition TÜV Rheinland Certificate No.

IEC 60950-1:2001, 1st Edition CB Scheme Certificate No.

Evaluated to all CB Countries

UL 60950-1, 1st Edition, CSA C22.2 No. 60950-1-03 File: Vol. Sec.

Supplementary Information: This product was tested and complies with all the requirements for the CE Mark.

/S/ _____
Dennis P. Symanski
Manager, Compliance Engineering
Sun Microsystems, Inc.
4150 Network Circle, MPK15-102
Santa Clara, CA 95054 U.S.A.
Tel: 650-786-3255
Fax: 650-786-3723

DATE

/S/ _____
Donald Cameron
Program Manager
Sun Microsystems Scotland, Limited
Blackness Road, Phase I, Main Bldg
Springfield, EH49 7LR
Scotland, United Kingdom
Tel: +44 1 506 672 395 Fax: +44 1 506 670 011

DATE

Preface

The Sun™ Dual Port 4x IB Host Channel Adapter PCI-X Installation and User's Guide provides installation instructions for the Sun Dual Port 4x IB Host Channel Adapter PCI-X card. This manual includes an overview of the adapter, supported platforms, cable information, and related documentation.

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

Note – The Sun Dual Port 4x IB Host Channel Adapter PCI-X card is an InfiniBand (IB) host channel adapter (HCA) card. The adapter is referred to as *IB-HCA* in this manual.

How This Book Is Organized

The manual is organized as follows:

[Chapter 1](#) provides a description of the adapter.

[Chapter 2](#) details how to install the adapter card and cables, and how to verify the installation.

[Chapter 3](#) contains information for configuring Internet Protocol over InfiniBand (IPoIB).

[Chapter 4](#) gives an overview of the InfiniBand software.

[Appendix A](#) contains information on the InfiniBand cable and connector pin outs.

[Appendix B](#) provides an overview of the SunVTS™ diagnostic software and the `ibhctest` test. The appendix also includes some troubleshooting tasks.

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 ¹	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. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
AaBbCc123	Book titles, new words or terms, words to be emphasized	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.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

¹ The settings on your browser might differ from these settings.

Related Documentation

The following documents provide additional information regarding the IB-HCA and related products.

Sun Dual Port 4x IB Host Channel Adapter PCI-X Release Notes (819-1281)

System Administration Guide: Devices and Files System (817-5093)

System Administration Guide: Network Services (816-4555)

Sun InfiniBand Switch Nine-Port Hardware Installation Guide (819-0504)

Sun InfiniBand Switch Nine-Port Administration Guide (819-0502)

Sun InfiniBand Switch Nine-Port Release Notes (819-1278)

PCI-X Addendum to the PCI Local Bus Specification, Revision 2.0

Documentation, Support, and Training

Sun Function	URL	Description
Documentation	http://www.sun.com/documentation/	Download PDF and HTML documents, and order printed documents
Support and Training	http://www.sun.com/supporttraining/	Obtain technical support, download patches, and learn about Sun courses

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.

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

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 Dual Port 4x IB Host Channel Adapter PCI-X Installation and User's Guide, part number 819-1280-10

Product Overview

This chapter provides an overview of the Sun Dual Port 4x IB Host Channel Adapter PCI-X card and includes the following sections:

- [“Product Features” on page 1](#)
- [“Hardware Description” on page 1](#)
- [“I/O Interfaces” on page 3](#)
- [“LED Displays” on page 3](#)
- [“Hardware, Software, and Cable Requirements” on page 5](#)

Product Features

Sun Dual Port 4x IB Host Channel Adapter PCI-X card features include:

- Dual 10 Gb/sec (4x) InfiniBand ports
- Integrated serializer/deserializer (SerDes)
- PCI-X Revision 2.2-compatible
- 64-bit or 32-bit 133 MHz PCI-X
- InfiniBand Trade Association (IBTA) version 1.1 interoperability
- Dual copper connectors (MicroGigaCN)

Hardware Description

The Sun Dual Port 4x IB Host Channel Adapter PCI-X card is an InfiniBand (IB) host channel adapter (HCA) card and is referred to as the IB-HCA in this manual. This adapter is a PCI-X Revision 2.2, 64-bit, 133MHz compatible card with dual 4x MicroGigaCN InfiniBand-compliant connectors for copper cables. [FIGURE 1-1](#) shows the IB-HCA card.

The board characteristics are:

- Power: approx. 12 watts total
- Low profile PCI form-factor card (approx. 2.5 in. by approx. 6.7 in.)
- Requires 4x InfiniBand cables
- 3.3V PCI-X or PCI slot (Note that the IB-HCA can not be plugged into a 5V connector.)
- Link status indication LED
- Two 4x InfiniBand (IB) ports
- Two faceplate sizes: standard and low profile

The standard PCI/PCI-X interface is used for data traffic, management traffic, and to generate and receive inbound configuration packets.




FIGURE 1-1 Sun Dual Port 4x IB Host Channel Adapter PCI-X Card

I/O Interfaces

The IB-HCA includes the following interfaces:

- Two 4x InfiniBand copper connectors
- PCI-X 64-bit 133 MHz edge connector

InfiniBand Interface

The IB-HCA provides two 4x InfiniBand v1.1 connector (or ports) for external copper cables and is IBTA specification 1.1-compliant.

PCI-X Interface

The PCI-X bus is a PCI-X version 2.2-compliant 64-bit 133 MHz interface. The adapter can be either a master initiating the PCI bus operations or a slave responding to PCI bus operations. The PCI/PCI-X bus can connect to a host CPU in an HCA application.

LED Displays

Two InfiniBand ports connectors are located on the front panel of the adapter. Each port has two LEDs, as shown in [FIGURE 1-2](#). An explanation of the LEDs is given in [TABLE 1-1](#).

TABLE 1-1 IB-HCA LEDs

Color	Meaning
Green	Physical Link – The physical link LED illuminates once the Tavor driver is successfully attached to the port.
Amber	Logical Link – The logical link LED illuminates once the InfiniBand port is in the Up state. This LED also blinks when there is traffic on the port.




FIGURE 1-2 LED Layout

Hardware, Software, and Cable Requirements

The hardware, software, and cable requirements for the Sun Dual Port 4x IB Host Channel Adapter PCI-X card are listed in [TABLE 1-2](#).

TABLE 1-2 Hardware, Software, and Cable Requirements

Operating System	Servers
Solaris 10 SPARC	Sun Fire™ V240, V440, V490, V890, V1280, E2900, E4900, E6900, E20K, E25K
Solaris 10 X86	Sun Fire V20z, V40z
InfiniBand Switch	
X3152A	Sun™ InfiniBand Switch P9
Diagnostics Version	
	SunVTS Version 6.0 Patch Set 1 and subsequent compatible releases
Supported Cables	
X9280A	Sun 4x-to-12x InfiniBand cable, 2 meter length
X9281A	Sun 4x-to-12x InfiniBand cable, 5 meter length
Note: Use only supported cables to ensure reliable InfiniBand interface connections. These cables can be ordered from Sun Microsystems, Inc.	

Adapter Installation

This chapter describes how to install the adapter in your system and verify that it has been installed correctly.

IB-HCA Bus and Slot Selection

For optimal performance, the IB-HCA card should be installed in the fastest available bus, preferably a 133 MHz PCI-X bus. But the IB-HCA can operate on the slower PCI-X and PCI buses. The IB-HCA should also be installed on its own PCI leaf. Plugging the IB-HCA card into a PCI leaf shared by another (slower) device can result in lower performance.

▼ Installing the Adapter



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 anti-static grounding straps and antistatic mats to help prevent damage due to electrostatic discharge.

Note – Refer to your system installation or service manual for detailed instructions for the following steps.

1. **Power off your server, using the standard shut down procedures described in the *Solaris Handbook for Sun Peripherals* or your system service manual.**

The *Solaris Handbook for Sun Peripherals* is shipped with the Solaris Operating System software and 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 3.3 Volt PCI-X or PCI 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.

Note – On the Sun Fire V20z, remove the top cover when connecting or disconnecting the InfiniBand I/O cable.

7. **Replace the cover on the unit.**
8. **If not already connected, connect the 12x end of the InfiniBand I/O cable(s) to the appropriate ports on the InfiniBand switch.**
See [FIGURE 2-1](#) for a typical adapter to switch cabling example.
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 [“Verifying the Installation”](#) on page 9.




FIGURE 2-1 Typical Adapter-to-Switch Configuration

▼ Verifying the Installation

Before you can verify the installation, you must install the adapter in the chassis, power the server, and cable it to an operational InfiniBand switch. The InfiniBand switch should automatically recognize InfiniBand servers when they are connected to the fabric.

- 1. 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. Also check if the Tavor driver is attached to the port. See [“Troubleshooting Tasks”](#) on page 25 for details.

2. 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 B, “Diagnostic Software and Troubleshooting”](#) on page 23 for more information.

Configuring IP Over InfiniBand

This chapter contains the following sections:

- [“IPoIB Configuration” on page 11](#)
- [“Managing IB Devices for IPoIB” on page 13](#)

IPoIB Configuration

IP Over InfiniBand (IPoIB) is automatically configured when the system is booted. The following procedure can be used when making changes to your configuration. See *System Administration Guide: Basic Administration (817-1985)* for Solaris setup information.

▼ Configuring IPoIB

1. **Confirm that the IB-HCA ports are up.**

To do this, type the command:

```
# dmesg | grep tavor
```

The output shows recently printed 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:

```
Dec 10 16:41:05 hostname pcplusmp: [ID 637496 kern.info] pcplusmp: pci15b3,5a44
(tavor) instance 0 vector 0x1c ioapic 0x4 intin 0x0 is bound to cpu 1
Dec 10 16:41:05 hostname pci_pci: [ID 370704 kern.info] PCI-device:
pci15b3,5a44@0, tavor0
Dec 10 16:41:05 hostname genunix: [ID 936769 kern.info] tavor0 is
/pci@0,0/pci1022,7450b/pci15b3,5a46@1/pci15b3,5a44@0
Dec 10 16:41:05 hostname tavor: [ID 548581 kern.info] tavor0: FW ver:
0003.0002.0000, HW rev: a1
Dec 10 16:41:05 hostname tavor: [ID 827127 kern.info] tavor0: MT23108
InfiniHost Mellanox Technologies (0x0002c9010976bb40)
Dec 10 16:41:06 hostname tavor: [ID 979957 kern.info] NOTICE: tavor0:
port 1 up (link width: 4x)
```

Refer to `dmesg(1M)` for more information.

2. Create the InfiniBand devices using the following command:

```
# devfsadm -c
```

A `/dev/ibdn` device is created for each IB-HCA port in the system where *n* is the port number. Refer to `devfsadm(1M)` for more information.

3. Verify InfiniBand devices.

Check that the InfiniBand devices were created. To do this, type the following command:

```
# ls -l /dev/ibd*
```

A list of InfiniBand devices (`/dev/ibd*`) is displayed.

4. Plumb each of the IB devices.

To do this, type the following command for each IB-HCA port in the system:

```
# ifconfig ibdn plumb host-ipv4-address up
```

where *n* is the instance number of the interface and `host-ipv4-address` is the designated IPv4 address for the device.

If you are using IPv6, use the following command instead:

```
# ifconfig ibdn inet6 plumb host-ipv6-address up
```

5. Verify that the devices are plumbed.

To do this, type the command:

```
# ifconfig -a
```

The output includes information on the IB devices.

6. Check that data can be send over the InfiniBand link.

Use the `ping` command to send ICMP (ICMP6) ECHO_REQUEST packets and receive ICMP (ICMP6) ECHO_RESPONSE packets between InfiniBand hosts. For example:

```
# ping host-ip-address
```

where *host-ip-address* is the IP address of another IB host. See `ping(1M)` for details. Repeat this procedure on all hosts that are connected to the IB switch.

Managing IB Devices for IPoIB

For information on managing the IB devices, see the *System Administration Guide: Devices and Files System* (817-5093).

InfiniBand Software Overview

This chapter includes the following topics:

- [“InfiniBand Software”](#) on page 15
- [“Related Software Documentation”](#) on page 18

InfiniBand Software

The InfiniBand software is included in the Solaris 10 Operating System. [FIGURE 4-1](#) shows a block diagram of the InfiniBand software stack which includes the following modules:

- Solaris 10 SPARC® Platform Edition or Solaris 10 X86 Platform Operating Systems
- [User Direct Access Programming Library \(uDAPL\)](#)
- [IP over InfiniBand Driver \(IPoIB\)](#)
- [InfiniBand Transport Framework \(IBTF\)](#)
- [InfiniBand Management Framework \(IBMF\)](#)
- [Tavor HCA Driver](#)




FIGURE 4-1 InfiniBand Software Stack

User Direct Access Programming Library

This library supports service providers written according to the uDAPL 1.2 specification. Sun provides default service provider modules to support the Tavor HCA driver.

IP over InfiniBand Driver

The IPoIB driver implements the IETF IPoIB protocol and provides IPoIB services for all IB-HCA ports present in the system. This driver enables the standard Solaris TCP/IP stack to operate over Infiniband. The link layer driver encapsulates IPv4/6 packets, adds IB multicast support, extends ARP/ND address resolution, and changes the network stack to support link addresses greater than 6 bytes.

InfiniBand Transport Framework

IBTF is the Solaris implementation of the InfiniBand Transport Layer (IBTL). IBTF consists of InfiniBand Device Manager (IBDM), InfiniBand Communications Manager (IBCM) and an InfiniBand Nexus driver. The Nexus driver along with IBDM creates and manages InfiniBand devices. The following devices are supported:

- IB I/O controller (IOC) devices
- IB pseudo devices
- IB service devices
- HCA, port, virtual physical point of attachment (VPPA) service devices

See `ib(4)` and `ib(7d)` for details.

IBCM provides a transport layer abstraction to IB clients to set up reliable connected channels along with service, multicast, and path lookup-related functionality. See `ibcm(7d)` for details.

InfiniBand Management Framework

IBMF is a complementary component to IBTF that enables IB clients to provide and access management services over the IB fabrics. It has two logical components. One component deals with sending and receiving management data to any type of management entity. The other component deals with sending and receiving management data to the Subnet Administrator (SA) only. Since every IB client needs to send and receive management data to the SA, IBMF provides functions that specifically handle the SA interaction.

Tavor HCA Driver

The Tavor HCA driver controls the Tavor hardware and enables the transfer of data over the IB wire. This driver interfaces with the IBTF and provides an implementation of Channel Interfaces that are defined by IBTF. It also enables management applications and agents to access the IB fabric.

Related Software Documentation

Manuals

- *System Administration Guide: Devices and Files System* (817-5093)
For information on InfiniBand devices, see the *System Administration Guide: Devices and Files System* (817-5093)
- *System Administration Guide: Network Services* (816-4555)
For general information on the InfiniBand network, see the *System Administration Guide: Network Services* (816-4555).
- *Sun Cluster 3.x Hardware Administration Manual for Solaris OS*
For information on administering Sun Cluster hardware, see the *Sun Cluster 3.x Hardware Administration Manual for Solaris OS* (817-0168)

Note – You can view or print Sun documentation at the following web site:
<http://www.sun.com/documentation>

Man Pages

- `cfgadm_ib(1M)` – InfiniBand hardware-specific commands for `cfgadm`
- `datadm(1M)` – maintain DAT static registry file
- `ifconfig(1M)` – configure network interface parameters
- `libdat(3LIB)` - direct access transport library
- `ib(4)` – InfiniBand device driver configuration files
- `ibmf(7)` – InfiniBand Management Transport Framework
- `daplt(7D)` – Tavor uDAPL service driver
- `ib(7D)` – InfiniBand Bus Nexus Driver
- `ibcm(7D)` – InfiniBand Communication Manager
- `ibd(7D)` – Infiniband IPoIB device driver
- `ibdm(7D)` – Solaris InfiniBand device manager
- `tavor(7D)` – InfiniBand (IB) Tavor driver

InfiniBand Cable Pin Assignments

FIGURE A-1 shows the InfiniBand cable connectors and pin assignments.




FIGURE A-1 Sun 4x-to-12x InfiniBand Cable Pin Assignments

Diagnostic Software and Troubleshooting

This appendix provides an overview of the SunVTS™ diagnostic application and troubleshooting tips for checking the adapter. There is also a section outlining some common troubleshooting issues. This appendix contains the following sections:

- “SunVTS Diagnostic Testing” on page 23
- “Troubleshooting Tasks” on page 25
- “Other Useful Utilities” on page 26

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 not currently available for Solaris x86 Operating Systems.

Refer to the SunVTS documents (listed in [TABLE B-1](#)) for instructions on how to run and monitor the `nettest` diagnostic. These SunVTS documents are available online at the following URL:

http://www.sun.com/products-n-solutions/hardware/docs/Software/system_administration/tools/testing_valid/sunvts/index.html

Select the document for the Solaris release on your system

TABLE B-1 SunVTS Documentation

Title	Descriptions
<i>SunVTS 6.0 PS1 Documentation Supplement</i> (819-1804)	Describes the new SunVTS features and tests, including the <code>ibhctest</code> .
<i>SunVTS 6.0 User's Guide</i> (817-7664)	Describes the SunVTS diagnostic environment.
<i>SunVTS 6.0 Test Reference Manual</i> (817-7665)	Describes each SunVTS test and describes the various test options and command-line arguments.
<i>SunVTS 6.0 Quick Reference Card</i> (817-7686)	Provides an overview of the user interface.

Using the SunVTS `ibhctest`

The `ibhctest` diagnostic test checks the functionality of Sun Dual Port 4x IB Host Channel Adapter PCI-X 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`.

▼ Using the `ibhctest`

1. Ensure that the SunVTS software and the Tavor driver are installed on your system.

To do this, type the following command:

```
# 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 ibd $n$  down unplumb
```

where n is the instance number of the interface.

3. Refer to *SunVTS 6.0 PS1 Documentation Supplement* (819-1804) for instructions on how to run `ibhctest`.

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 package (`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  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

- adapter
 - description, 1
 - features, 1
 - installation, 7
 - layout, 2

C

- cable requirements, 5
- cables
 - connectors, 19
 - part numbers, 5
 - pin assignments, 19

D

- diagnostics, `ibhctest`, 24
- documentation
 - man pages, 18
 - related, xix, 18
 - SunVTS, 23

I

- I/O interfaces, 3
- IB-HCA
 - installation, 8
 - Sun Dual Port 4x IB Host Channel Adapter PCI-X card, 1
- `ibhctest` diagnostics, 24
- IBMF, 15, 17
- IBTF, 15, 17
- InfiniBand cable, 19
- InfiniBand Management Framework, 17

- InfiniBand software overview, 15
- InfiniBand Transport Layer, 17
- installation, adapter, 7
- interface
 - InfiniBand, 3
 - PCI-X, 3
- IPoIB, 15
 - configuring, 11
 - driver, 16
 - managing devices, 13

L

- LEDs
 - descriptions, 3
 - layout, 4
 - verification, 10

O

- operating systems, 5

R

- requirements
 - hardware, 5
 - software, 5

S

- servers, 5
- software packages, 25
- SunVTS
 - diagnostic application, 23
 - documentation, 23, 24
 - `ibhctest` diagnostics, 24

versions, 5

T

Tavor driver, 17

troubleshooting

 IB-HCA, 25

 tasks, 25

 utilities, 26

U

uDAPL, 15, 16

utilities, 26

 cfgadm, 26

 kstat, 27

 netstat, 27

 snoop, 27