

# **Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2**

User's Guide



Part No.: E23921-01  
April 2013

Copyright © 2010, 2013, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

---

Copyright © 2010, 2013, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

# Contents

---

## **Using This Documentation** v

### **Installing the Adapter** 1

- ▼ Install the Adapter 1
- Adapter Features 3
- Adapter Specifications 4
- InfiniBand Interface 5
- PCI Express Interface 5
- LEDs and Ports 6
- I<sup>2</sup>C Compatible Interface 7
- Node GUID 8
- Hardware and Software Requirements 8
- ▼ Hot-Swap the Adapter (Oracle Solaris) 8

### **Installing InfiniBand Software on the Oracle Solaris OS** 11

- InfiniBand Software for Oracle Solaris 10 11
- ▼ Download the Firmware Flash Update Tool for IB-HCAs (Oracle Solaris 10) 12
- ▼ Verify the Installation (Oracle Solaris) 13
- ▼ Update the Firmware (Oracle Solaris) 15
- InfiniBand Devices on the Oracle Solaris 10 OS 16

### **Installing InfiniBand Software on Linux** 17

- InfiniBand Software for Linux 17

- ▼ Acquire the BXOFED Software (Linux) 18
- ▼ Install the BXOFED Software (Linux) 20
- Internet Protocol Over InfiniBand (Linux) 21
- ▼ Verify the Installation (Linux) 21

**Replacing a Short Bracket With a Tall Bracket 23**

- ▼ Remove the Short Bracket From the Adapter 23
- ▼ Assemble and Install a Tall Bracket 25

**Glossary 27**

**Index 29**

# Using This Documentation

---

This guide provides an overview, installation instructions, and specifications of the Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2 from Oracle. The instructions in this guide are designed for system administrators with experience installing network hardware and software.

---

**Note** – This adapter is based on the Mellanox ConnectX 2 ASIC.

---

- “Product Notes” on page v
- “Related Documentation” on page vi
- “Feedback” on page vi
- “Access to Oracle Support” on page vi

---

## Product Notes

For late-breaking information and known issues about this product, refer to the product notes at:

<http://docs.oracle.com/cd/E19241-01/>

---

## Related Documentation

---

Documentation	Links
All Oracle products	<a href="http://docs.oracle.com">http://docs.oracle.com</a>
Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2	<a href="http://docs.oracle.com/cd/E19241-01/">http://docs.oracle.com/cd/E19241-01/</a>
Oracle Integrated Lights Out Manager (ILOM)	<a href="http://www.oracle.com/goto/ILOM/docs">http://www.oracle.com/goto/ILOM/docs</a>
Oracle Solaris 11 OS	<a href="http://www.oracle.com/goto/Solaris11/docs">http://www.oracle.com/goto/Solaris11/docs</a>
Oracle Solaris 10 OS	<a href="http://www.oracle.com/goto/Solaris10/docs">http://www.oracle.com/goto/Solaris10/docs</a>
Oracle VM Server for SPARC	<a href="http://www.oracle.com/goto/VM-SPARC/docs">http://www.oracle.com/goto/VM-SPARC/docs</a>
Oracle VTS	<a href="http://www.oracle.com/goto/VTS/docs">http://www.oracle.com/goto/VTS/docs</a>

---

---

## Feedback

Provide feedback about this documentation at:

<http://www.oracle.com/goto/docfeedback>

---

## Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

# Installing the Adapter

---

The adapter is a dual-port 4x InfiniBand PCIe Generation 2 low-profile adapter.

- [“Adapter Features” on page 3](#)
- [“Adapter Specifications” on page 4](#)
- [“InfiniBand Interface” on page 5](#)
- [“PCI Express Interface” on page 5](#)
- [“LEDs and Ports” on page 6](#)
- [“I2C Compatible Interface” on page 7](#)
- [“Node GUID” on page 8](#)
- [“Hardware and Software Requirements” on page 8](#)
- [“Install the Adapter” on page 1](#)
- [“Hot-Swap the Adapter \(Oracle Solaris\)” on page 8](#)

## Related Information

- [“Installing InfiniBand Software on the Oracle Solaris OS” on page 11](#)
- [“Installing InfiniBand Software on Linux” on page 17](#)

---

## ▼ Install the Adapter

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



---

**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 adapter in a system, use antistatic grounding straps and antistatic mats to help prevent damage due to electrostatic discharge.

---

**1. Check that the bracket on the adapter is the correct size for your system.**

An alternative tall bracket is supplied with the adapter. If you need to use a different bracket, perform the instructions in [“Replacing a Short Bracket With a Tall Bracket”](#) on page 23.

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

Product documentation for Sun servers is available at <http://docs.sun.com>.

**3. Remove the cover from the system to access the adapter slots and connectors.**

**4. Select an available PCIe x8 slot and remove the blank filler panel for that slot.**

Or, if you are replacing an existing adapter in that slot, remove the adapter.

**5. Install the adapter into the slot, pushing the adapter’s edge connector into the connector on the chassis.**

Ensure that the front plate on the adapter mounts flush with the chassis panel opening.

**6. If applicable, install the screw in the front plate to secure the adapter into the chassis.**

**7. Attach the 4x end of each InfiniBand I/O cable to an IB-HCA port connector.**

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.

---

**8. Replace the cover on the unit.**

**9. If not already connected, connect the other 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.

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

This step completes the hardware installation.

**11. Verify the installation.**

See [“Verify the Installation \(Oracle Solaris\)”](#) on page 13 or [“Verify the Installation \(Linux\)”](#) on page 21.

**Related Information**

- [“Installing InfiniBand Software on the Oracle Solaris OS”](#) on page 11
- [“Installing InfiniBand Software on Linux”](#) on page 17

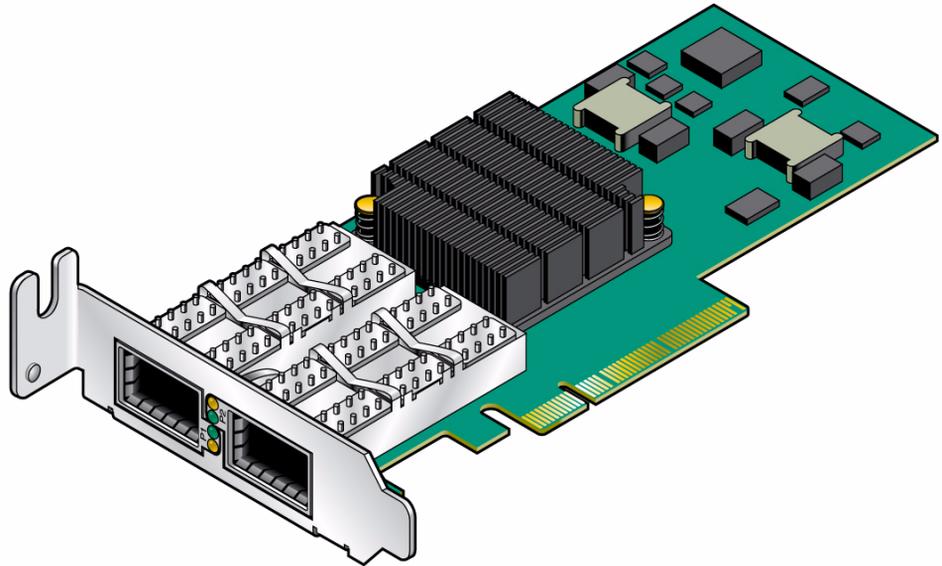
---

# Adapter Features

The main features of the adapter are as follows:

- PCIe 2.0 adapter
- PCI Express expansion board with an x8 edge connector compatible to the PCI Express 1.0a specification
- Two 4x InfiniBand copper ports for connecting InfiniBand traffic (4x IB connectors)
- 4x IB port speed support: 10 Gbps, 20 Gbps, or 40 Gbps (QDR - Quad data rate)
- Short mounting bracket (an alternative tall bracket is supplied along with the adapter)
- IBTA v1.2 compliant
- Media detect circuit, which supports external InfiniBand fiber solutions
- European Union Restriction of Hazardous Substances (RoHS) compliant

**FIGURE:** Adapter With Short Bracket



## Related Information

- [“Adapter Specifications” on page 4](#)

---

# Adapter Specifications

The adapter receives power from the PCI Express edge connector. All other required power voltages are generated by on-board switch mode regulators.

**TABLE:** Specifications

---

<b>Physical</b>	
Size	2.54 in. x 5.37 in. (64.4 mm x 136.47 cm)
Air flow	200 LFM @ 55°C
4x 10Gbps connector	InfiniBand (copper, current rating: 0.5A max) with active media adapter support

---

<b>Protocol Support</b>	
InfiniBand	IBTA v1.2.1, auto-negotiation 40 Gbps, 5 Gbps 20 Gbps, 5 Gbps 10 Gbps, 2.5 Gbps
QoS	8 InfiniBand virtual lanes for each port
RDMA support	All ports
PCI Express	2.0 SERDES @ 5.0 GT/s

---

<b>Power and Environmental</b>	
Voltage	12 V, 3.3 V
Typical power	12.23 W
Maximum power	14.32 W
Temperature	0°C to 55°C

---

<b>Regulatory</b>	
Safety	IEC/EN 60950-1:2001, ETSI EN 300 019-2-2
Environmental	IEC 60068-2- 64, 29, 32
RoHS	RoHS-R5

---

## Related Information

- [“Adapter Features” on page 3](#)

---

# InfiniBand Interface

The adapter is compliant with the *InfiniBand Architecture Specification, Release 1.2*. The adapter has two compliant 4x InfiniBand ports, 1 and 2. The adapter provides access to these ports by means of two 4x InfiniBand [QSFP](#) connectors for external InfiniBand copper cables. These cables must be compliant with the *InfiniBand Architecture Specification, Release 1.2*. Connector 1 connects to Port 1 of the device, while connector 2 connects to Port 2.

The adapter is embedded with a media detect circuit, which supports external InfiniBand fiber solutions. These external devices are connected to the InfiniBand port connectors using active media converters, such as the Emcore QTR3400 Smart Module or the Fujitsu FPD-010R008-0E o-microGiGaCN.

## Related Information

- [“PCI Express Interface”](#) on page 5
- [“LEDs and Ports”](#) on page 6
- [“I2C Compatible Interface”](#) on page 7

---

# PCI Express Interface

The adapter has eight Tx/Rx pairs of [SerDes](#) providing for a PCI Express x8 edge connector interface, version 1.0a compatible. The adapter can be either a master, initiating the PCI Express bus operations, or a slave responding to PCI bus operations.

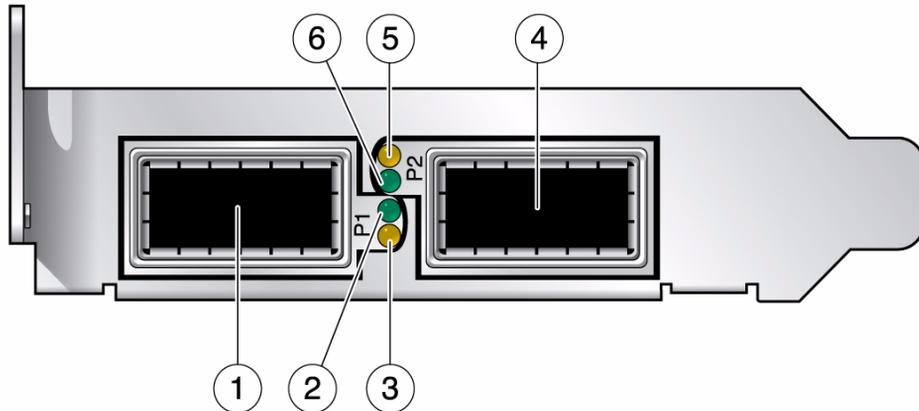
## Related Information

- [“InfiniBand Interface”](#) on page 5
- [“LEDs and Ports”](#) on page 6
- [“I2C Compatible Interface”](#) on page 7

# LEDs and Ports

The adapter has four LEDs located on the I/O panel. Two LEDs are assigned to each 4X port.

**FIGURE:** I/O Panel With Dual Ports and LEDs



## Figure Legend

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

The same port names and LED footprints apply when a tall bracket is installed on the adapter. To install the supplied tall bracket, see [“Replacing a Short Bracket With a Tall Bracket”](#) on page 23.

The pair of LEDs for each port have the meanings described in the following table.

**TABLE:** LED Meanings

LED Color	LED Name	LED State	Meaning
Green	Physical Link	Lit	The link bringup process has successfully completed and the link width, link speed, link polarity, and link reversal have been negotiated with the neighbor port on the other end of the cable.
		Unlit	A physical connection has not been established.
Amber	Data Activity	Steady light	Infiniband is discovered over the physical link, but no data is being passed.
		Blinking light	Data is being passed.
		Unlit	A physical or logical connection (or both) has not been established.

### Related Information

- [“InfiniBand Interface” on page 5](#)
- [“PCI Express Interface” on page 5](#)
- [“I2C Compatible Interface” on page 7](#)

---

## I<sup>2</sup>C Compatible Interface

A three-pin header, designated with reference name J5 on the adapter, is provided as the I<sup>2</sup>C compatible interface.

### Related Information

- [“InfiniBand Interface” on page 5](#)
- [“PCI Express Interface” on page 5](#)
- [“LEDs and Ports” on page 6](#)

---

# Node GUID

A label on the back of the adapter displays the 64-bit GUID. This GUID uniquely identifies this adapter in the server and on the IB fabric.

## Related Information

- “Verify the Installation (Oracle Solaris)” on page 13
- “Update the Firmware (Oracle Solaris)” on page 15

---

# Hardware and Software Requirements

For the latest list of supported platforms and operating systems, refer to the *Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2 Product Notes*. This document is available at: <http://docs.oracle.com/cd/E19241-01/>.

## Related Information

- “Install the Adapter” on page 1

---

## ▼ Hot-Swap the Adapter (Oracle Solaris)

---

**Note** – Hot-swapping is not supported with Oracle Solaris 10 10/09. Hot-swapping is supported beginning with Oracle Solaris 10 9/10. Currently, hot-swapping is supported on the following servers: Sun SPARC Enterprise M4000, M5000, M8000, and M9000 Servers.

---

This task is for removing and installing an adapter without powering off the system.

1. Add this line to the `/etc/system` file:

```
set pcie:pcie_disable_ari = 1
```

Reboot the system after the modification for the changes to take effect.

2. Unplumb the adapter:

```
# ifconfig ibd(0/1 or 2/3 ...) unplumb
```

3. Find the device number as required for hot-swapping:

---

**Note** – The `Ap_Id` changes from platform to platform. Use the `Ap_Id` related to the `ib/hp` type. Refer to the hot-plug/hot-swap procedures in the Solaris 10 documentation collections at: <http://docs.sun.com>

---

The following is an example. The output from your platform could look different.

```
# cfgadm
Ap_Id          Type          Receptacle  Occupant    Condition
...
PCI-LP0        ib/hp         connected   configured  ok
PCI-LP1        ib/hp         connected   configured  ok
c2             scsi-sas      connected   unconfigured unknown
c3             scsi-sas      connected   unconfigured unknown
c4             scsi-sas      connected   unconfigured unknown
hca:2C90109763F70  IB-HCA       connected   configured  ok
hca:2C90109763F71  IB-HCA       connected   configured  ok
ib             IB-Fabric    connected   configured  ok
usb0/1        unknown      empty       unconfigured ok
```

4. Unconfigure the adapter:

```
# cfgadm -c unconfigure PCI-LP0
```

5. Disconnect the adapter:

```
# cfgadm -c disconnect PCI-LP0
```

6. Remove the adapter from the server.

7. Install a new adapter.

See “Install the Adapter” on page 1.

8. Connect the adapter.

```
# cfgadm -c connect PCI-LP0
```

9. Configure the adapter.

```
# cfgadm -c configure PCI-LP0
```

10. Plumb the adapter.

```
# ifconfig ibdn plumb IP_address up
```

11. Verify the installation.

See “Verify the Installation (Oracle Solaris)” on page 13.

**Related Information**

- [“Install the Adapter” on page 1](#)

# Installing InfiniBand Software on the Oracle Solaris OS

---

These topics provide an overview of installing and using the InfiniBand software stack for the Oracle Solaris OS.

Solaris 11 is supported. Consult the product notes for your server for recent information about supported operating systems, firmware and software updates, and other issues not covered in the main product documentation.

- [“InfiniBand Software for Oracle Solaris 10”](#) on page 11
- [“Download the Firmware Flash Update Tool for IB-HCAs \(Oracle Solaris 10\)”](#) on page 12
- [“Verify the Installation \(Oracle Solaris\)”](#) on page 13
- [“Update the Firmware \(Oracle Solaris\)”](#) on page 15
- [“InfiniBand Devices on the Oracle Solaris 10 OS”](#) on page 16

## Related Information

- [“Installing InfiniBand Software on Linux”](#) on page 17

---

## InfiniBand Software for Oracle Solaris 10

InfiniBand is a network architecture for the large-scale interconnection of computing and I/O nodes through a high-speed switched fabric. To operate InfiniBand on a Sun server, you need an InfiniBand HCA (the adapter) and an InfiniBand software stack.

InfiniBand software is bundled with the Oracle Solaris 10 OS. The package containing the device driver for the adapter is `SUNWhermon`. The driver name is `hermon`.

For details about InfiniBand software supported in Oracle Solaris 10 releases, refer to the following documents in the Solaris 10 Release and Installation Collection available at <http://docs.sun.com>:

- *Solaris 10 What's New*
- *Solaris 10 Release Notes*
- *Solaris 10 Package List*

---

**Note** – You must use the SUNwhermon package that is available in the Oracle Solaris 10 10/09 OS and subsequent Oracle Solaris releases with this adapter.

---

The InfiniBand software stack, consisting of the upper-layer protocols and transport framework, is included in all of the Solaris software groups described in the *Solaris Installation Guide*. The SUNwhermon package is included in the Entire+OEM, Entire, and Developer software groups. If you are not using any of these groups, you must explicitly add the SUNwhermon package during initial installation. If you are not installing software, use the `pkgadd(1)` utility to add the package prior to using the adapter.

#### **Related Information**

- “Download the Firmware Flash Update Tool for IB-HCAs (Oracle Solaris 10)” on page 12
- “Verify the Installation (Oracle Solaris)” on page 13
- “Update the Firmware (Oracle Solaris)” on page 15
- “InfiniBand Devices on the Oracle Solaris 10 OS” on page 16

---

## ▼ Download the Firmware Flash Update Tool for IB-HCAs (Oracle Solaris 10)

---

**Note** – The firmware flash utility bundled with Oracle Solaris 11 11/11 and later releases is available as the `fwflash(1M)` command. The `fwflash` command supports the adapter.

---

The `firmwareflash` utility in the Oracle Solaris 10 OS does not support the adapter. You must download a separate package (version v2.0 or higher) containing that utility from the Oracle Download web page at:

<http://www.oracle.com/technology/software/index.html>

#### **1. Go to the Download A-Z tab and search for “Sun Firmware Flash Utility.”**

Refer to the installation instructions in the package README file.

## 2. Check that the correct version is installed:

```
# firmwareflash -v  
firmwareflash: version v2.0
```

---

**Note** – This command must display version number 2.0 or higher.

---

### Related Information

- [“Verify the Installation \(Oracle Solaris\)”](#) on page 13
- [“Update the Firmware \(Oracle Solaris\)”](#) on page 15
- [“InfiniBand Devices on the Oracle Solaris 10 OS”](#) on page 16

---

## ▼ Verify the Installation (Oracle Solaris)

1. **Install the adapter in the chassis.**  
See [“Install the Adapter”](#) on page 1.
2. **Power on the server and cable the server to an operational InfiniBand switch.**  
Refer to the documentation for your server.
3. **Ensure that the cables are connected to the adapter and switches.**
4. **Verify that the IB Subnet Manager is in operation on the network.**  
Refer to the documentation for your network hardware for more information.
5. **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.
6. **Check that the amber LED is illuminated for each port that is connected to the switch.**
7. **Verify that the IB-HCA ports are up and the driver is attached.**

**a. Obtain the state of the device installed.**

```
# cfgadm -als "cols=ap_id:condition" hca
Ap_Id                               Condition
hca:2C90109763F70                   ok
```

If more than one IB-HCA device is installed in the server, a row is displayed for each. Look for the row displaying `hca:GUID` where *GUID* is the 64-bit number from the physical label on the adapter. See [“Node GUID” on page 8](#).

The Condition column must display `ok` to indicate that the driver has discovered the hardware and is bound to it. Refer to the `cfgadm_ib(1m)` man pages for details about the IB specific extensions.

**b. Obtain port GUIDs for each port on the adapter.**

```
# cfgadm -als "cols=ap_id:info" hca
Ap_Id                               Information
hca:2C90109763F70                   VID: 0x15b3, PID: 0x5a44,
#ports: 0x2, port1 GUID: 0x2C90109763F71, port2 GUID:
0x2C90109763F72
```

If more than one IB-HCA device is installed in the server, a row is displayed for each device. Look for the row displaying `hca:GUID`, where *GUID* is the 64-bit number from the physical label on the adapter. See [“Node GUID” on page 8](#).

Use the port number and GUID displayed by this command for your IB-HCA device in the following step.

**c. Verify that the IB ports and partitions are configured by the Subnet Manager.**

```
# cfgadm -als "select=type(IB-VPPA),cols=ap_id"
Ap_Id
ib::2C90109763F71,ffff,ipib
ib::2C90109763F72,ffff,ipib
```

The command displays the `AP_ID` column: where each row has the format of `ib::Port GUID,P_Key,ipib`. Match the Port GUIDs from the previous command with these port GUIDs. There must be one row corresponding to the port and *P\_Key* setup by the Subnet Manager. If an entry is missing, check the Subnet Manager configuration.

**Related Information**

- [“Download the Firmware Flash Update Tool for IB-HCAs \(Oracle Solaris 10\)” on page 12](#)
- [“Update the Firmware \(Oracle Solaris\)” on page 15](#)
- [“InfiniBand Devices on the Oracle Solaris 10 OS” on page 16](#)

---

## ▼ Update the Firmware (Oracle Solaris)

To use this adapter with the Solaris OS, the minimum firmware version must be 2.7.8130.

### 1. Display the revision level of your adapter.

```
# firmwareflash -l -c IB
```

Look for the revision number that appears after the `Firmware revision` string. If more than one IB-HCA device is displayed, look for the `Node Image GUID` that matches the GUID displayed on the physical GUID label of the adapter being installed. See “[Node GUID](#)” on page 8.

If the firmware version is not at 2.7.8130 or higher, you must update the firmware. Only update the firmware on your adapter with files specifically approved for the Sun product.

### 2. Select and download approved firmware files from:

<https://support.oracle.com/CSP/ui/flash.htm>

### 3. Use the `firmwareflash` command to install the firmware.

```
# firmwareflash -d device_path output -f firmware_bin_file
```

### 4. Reboot the system to enable the new firmware.

#### Related Information

- “[Download the Firmware Flash Update Tool for IB-HCAs \(Oracle Solaris 10\)](#)” on page 12
- “[Verify the Installation \(Oracle Solaris\)](#)” on page 13
- “[InfiniBand Devices on the Oracle Solaris 10 OS](#)” on page 16

---

# InfiniBand Devices on the Oracle Solaris 10 OS

For details about InfiniBand software stack configurations in a Oracle Solaris 10 Update release, refer to the *System Administration Guide: Devices and File Systems* document in the *Solaris 10 System Administrator Collection* available at <http://docs.sun.com>.

Section 9 of this guide, titled “Using InfiniBand Devices (Overview/Tasks),” describes how to set up upper layer protocols such as **IPoIB** and **uDAPL**.

When using IPoIB, verify that the `broadcast` group is configured by the Subnet Manager in the partition where the IPoIB link will be used. Refer to the documentation for the IB Subnet Manager for more information.

## Related Information

- “InfiniBand Software for Oracle Solaris 10” on page 11
- “Download the Firmware Flash Update Tool for IB-HCAs (Oracle Solaris 10)” on page 12
- “Verify the Installation (Oracle Solaris)” on page 13
- “Update the Firmware (Oracle Solaris)” on page 15

# Installing InfiniBand Software on Linux

---

These topics provide an overview of installing and using the InfiniBand software stack for Linux.

Consult the release notes for your server for recent information about supported operating systems, firmware and software updates, and other issues not covered in the main product documentation.

- [“InfiniBand Software for Linux” on page 17](#)
- [“Acquire the BXOFED Software \(Linux\)” on page 18](#)
- [“Install the BXOFED Software \(Linux\)” on page 20](#)
- [“Internet Protocol Over InfiniBand \(Linux\)” on page 21](#)
- [“Verify the Installation \(Linux\)” on page 21](#)

## Related Information

- [“Installing InfiniBand Software on the Oracle Solaris OS” on page 11](#)

---

## InfiniBand Software for Linux

InfiniBand is a network architecture for the large-scale interconnection of computing and I/O nodes through a high-speed switched fabric. To operate InfiniBand on an Oracle server, you need an InfiniBand HCA (the adapter) and an InfiniBand software stack.

---

**Note** – An InfiniBand software stack is part of the Oracle Linux distribution.

---

With most supported Linux releases, you must also install the OFED software stack (version 1.5.1 or higher). Refer to your Linux vendor for software installation recommendations and support.

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 (such as <http://OpenFabrics.org>) and your current vendors.

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

If you are running a Linux OS without OFED, or using the HCA in a network with the Sun Network QDR Infiniband Gateway Switch, you can download the BXOFED from My Oracle Support: <http://support.oracle.com>

### Related Information

- “Acquire the BXOFED Software (Linux)” on page 18
- “Install the BXOFED Software (Linux)” on page 20
- “Internet Protocol Over InfiniBand (Linux)” on page 21
- “Verify the Installation (Linux)” on page 21

---

## ▼ Acquire the BXOFED Software (Linux)

If you are running a Linux OS without OFED, or if you are using the adapter with the Sun Network QDR Infiniband Gateway Switch, download the BXOFED software stack from the My Oracle Support web site.

**1. Open a web browser on a host that will receive the BXOFED software.**

**2. Go to this URL:**

<http://support.oracle.com>

Oracle’s My Oracle Support page is displayed.

**3. Sign in if you already have an account.**

The dashboard page is displayed.

---

**Note** – If you do not have an account, you must register.

---

**4. Click the Patches & Updates tab.**

The Patches & Updates page is displayed.

5. **In the Patch Search window, click the click Product or Family (Advanced Search).**  
The Patch Search window updates.
6. **In the Product is field, type BridgeX.**  
Possible products are suggested.
7. **Click on the most appropriate link.**  
The Release is field might autopropagate with the most current version.
8. **In the Release drop-down menu, select the most current version of the BridgeX OFED software.**  
For example, BridgeX OFED 1.5.1.
9. **Click outside the drop-down menu.**
10. **Click Search.**  
The Patch Search window expands with the search results.
11. **In the Patch Name column, click the patch number link respective to your platform.**  
For example, 12621910. The Patch Search window reformats.
12. **Click Read Me to display the README file.**
13. **Click Download.**  
The File Download window opens.
14. **Click the *filename.zip* link to initiate the download.**  
For example, p12621910\_151\_Linux-x86-64.zip.
15. **Indicate where the file should be saved.**  
The file is downloaded and saved.
16. **In your receiving directory, decompress the *filename.zip* file.**  
The BXOFED software is in the BXOFED-1.5.1-*version\_for* Oracle.tgz file. There are also readme, release notes, installation guide and user manual files in the *filename.zip* file.
17. **Read the readme, release notes, and installation guide files for information how to install the BXOFED software.**

### **Related Information**

- [“Install the BXOFED Software \(Linux\)” on page 20](#)
- [“Verify the Installation \(Linux\)” on page 21](#)
- [“InfiniBand Software for Linux” on page 17](#)

---

## ▼ Install the BXOFED Software (Linux)

When you install the BXOFED software, any previous installations of OFED or BXOFED software are removed. Configuration files are not removed.

---

**Note** – If you are installing the BXOFED software on a cluster, install the software onto one of the cluster nodes, then install the `.rpm` files in the `OFED-1.5.1/RPMS` on all remaining cluster nodes using cluster-aware tools.

---

1. Become superuser of the host that received the BXOFED software.
2. Change to the directory where you extracted the `.tgz` file.
3. Run the installation script.

```
# ./BXOFED-1.5.1-1.6.3/install.pl
```

The script begins. Interactive menus direct you through the installation process. During the installation, two configuration files are created:

- `ofed.conf` – contains the names of the software modules installed and the configuration settings chosen during the installation.
- `ofed_net.conf` – contains the IPoIB configuration settings chosen during the installation.

The script finishes. This information is found in the respective locations:

- Man pages are installed in `/usr/share/man`.
- Documentation is installed under the `/usr/share/doc` directory.
- IPoIB configuration information is installed under the `/etc/sysconfig/network*` directory.
- The `openibd` daemon is installed under the `/etc/init.d` directory.
- BXOFED commands are located in the `/usr/bin` and `/usr/sbin` directories.
- BXOFED software installation information is displayed with the `/etc/infiniband/info` script.

4. (Optional) If the `.tgz` file was extracted to a NFS shared directory for a cluster, then to install the BXOFED software onto any remaining nodes in that cluster.
  - a. Log in as superuser of a node to receive the BXOFED software.
  - b. Change to the directory where the `.tgz` file was extracted.

c. Install the BXOFED software automatically.

```
# ./BXOFED-1.5.1-1.6.3/install.pl -c path/ofed.conf -n path/ofed_net.conf
```

where *path* is the directory path to the `ofed.conf` and `ofed_net.conf` files.

d. Repeat from [Step a](#) for all nodes to receive the BXOFED software.

5. Reboot the Linux InfiniBand host(s).

#### Related Information

- *Gateway Remote Administration*, upgrading the gateway firmware
- [“Acquire the BXOFED Software \(Linux\)”](#) on page 18
- [“Verify the Installation \(Linux\)”](#) on page 21

---

## Internet Protocol Over InfiniBand (Linux)

Support for [IPoIB](#) is included in the BXOFED software distribution.

See [“Acquire the BXOFED Software \(Linux\)”](#) on page 18.

#### Related Information

- [“InfiniBand Software for Linux”](#) on page 17
- [“Acquire the BXOFED Software \(Linux\)”](#) on page 18
- [“Install the BXOFED Software \(Linux\)”](#) on page 20
- [“Verify the Installation \(Linux\)”](#) on page 21

---

### ▼ Verify the Installation (Linux)

1. Install the adapter in the chassis.

See [“Install the Adapter”](#) on page 1.

**2. Power on the server and cable it to an operational InfiniBand switch.**

Refer to the documentation for your server.

---

**Note** – The InfiniBand switch should automatically recognize the adapter when it is connected to the fabric if the IB Subnet Manager is running on the switch or on a host within the subnet.

---

**3. Ensure that the cables are connected to the adapter and switches.**

**4. Verify that the IB Subnet Manager is running on the IB switch or on a host within the subnet.**

Refer to the manual for the IB Subnet Manager for more information.

**5. 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.

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

**7. Verify that the IB-HCA ports are up and the driver is attached:**

```
# ibstat
```

The output shows system diagnostic messages that have the string `m1x4` in the message (the name of the Linux driver). Included in the output is a message that indicates whether the port is up or down.

**Related Information**

- [“InfiniBand Software for Linux” on page 17](#)
- [“Acquire the BXOFED Software \(Linux\)” on page 18](#)
- [“Install the BXOFED Software \(Linux\)” on page 20](#)
- [“Internet Protocol Over InfiniBand \(Linux\)” on page 21](#)

# Replacing a Short Bracket With a Tall Bracket

---

By default, the Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2 from Oracle has a short bracket. A tall bracket is also provided.

- [“Remove the Short Bracket From the Adapter” on page 23](#)
- [“Assemble and Install a Tall Bracket” on page 25](#)

## **Related Information**

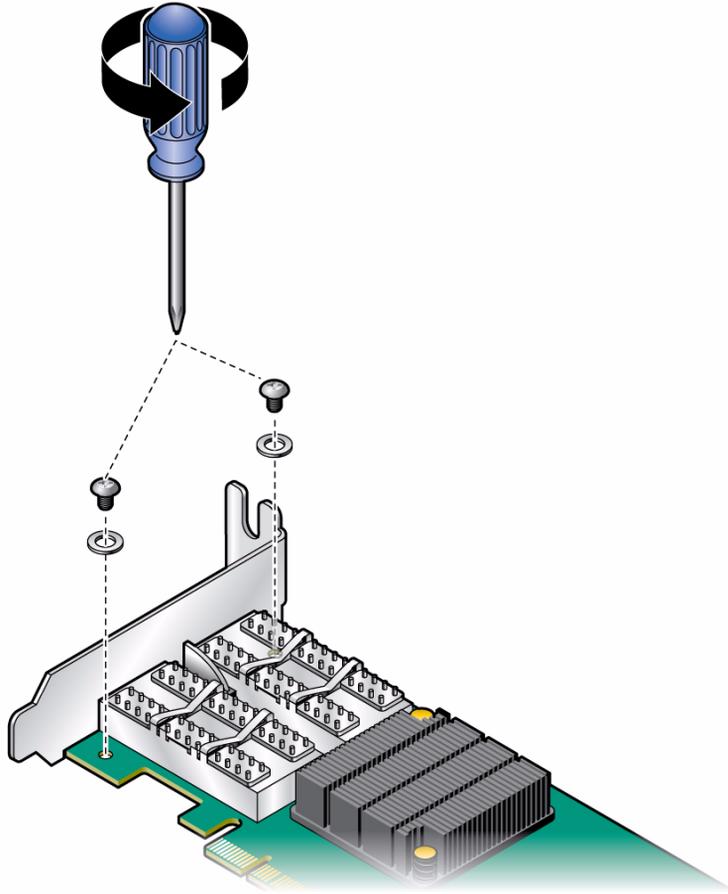
- [“Adapter Features” on page 3](#)

---

## ▼ Remove the Short Bracket From the Adapter

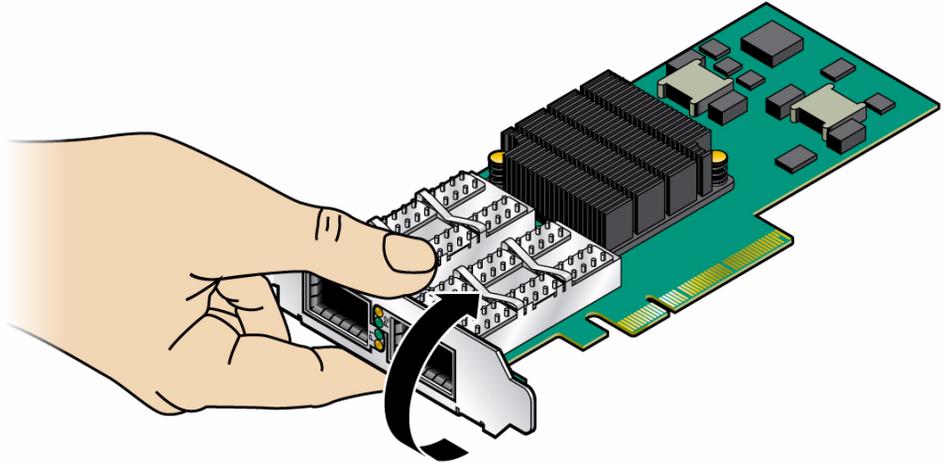
### **1. Remove the bracket screws and washers.**

Unscrew both screws from the adapter using a torque screwdriver.



**2. Detach the bracket.**

- a. Grip the bracket as shown in the following figure, placing your thumb on the LED component.**



- b. In a rotating move toward the component side of the adapter, slide the bracket out of the connectors.
- c. Gently hold your thumb on the LED component, as shown in the figure. At the same time extract the bracket, while making sure to protect the LEDs.
- d. Make a rotating move to detach the short bracket, as shown in the figure.

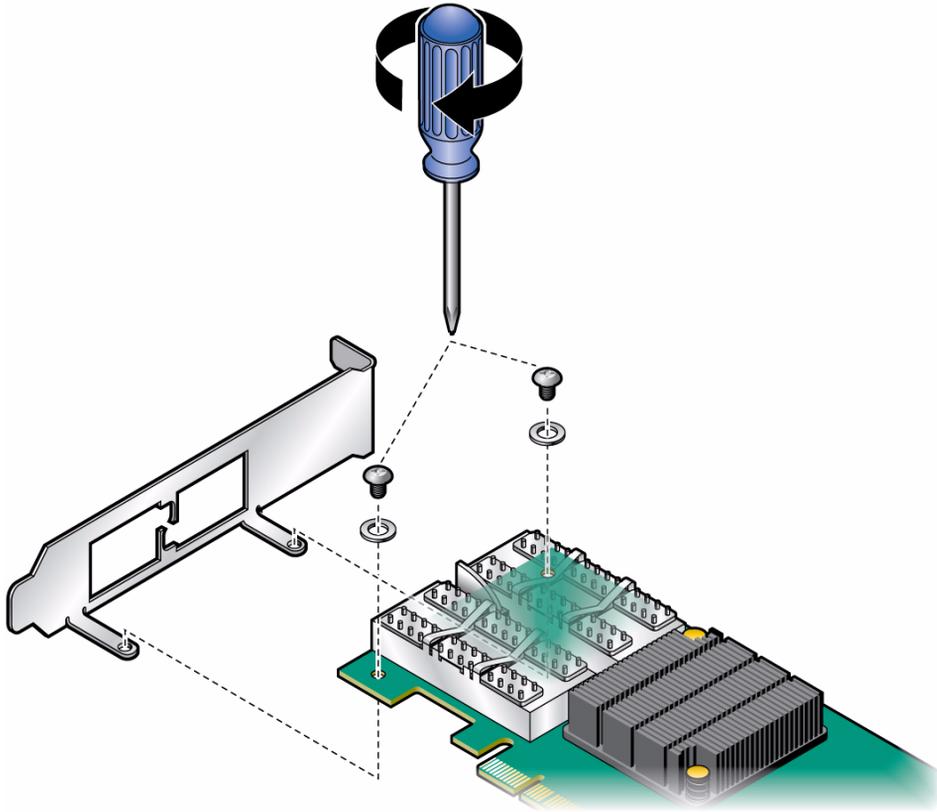
#### Related Information

- [“Assemble and Install a Tall Bracket” on page 25](#)

---

## ▼ Assemble and Install a Tall Bracket

1. Place the tall bracket onto the adapter.  
Gently fitting the connectors through the bracket connector holes.
2. Ensure that the LEDs are aligned into their intended bracket holes.



3. Insert a screw along with a washer into each of the two holes on the adapter intended for holding the bracket as shown in the figure.

Use a torque screwdriver to apply up to 2 lbs per inch torque on each screw.

#### **Related Information**

- [“Remove the Short Bracket From the Adapter”](#) on page 23

# Glossary

---

---

## A

**adapter** Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2.

---

## B

**BoIB** Boot over InfiniBand.

---

## I

**IB** InfiniBand. A switched fabric communications link primarily used in high-performance computing.

**IB-HCA** InfiniBand Host Channel Adapter.

**IBTA** InfiniBand Trade Association.

**IPoIB** Internet Protocol over InfiniBand.

---

## O

**OFED** OpenFabrics Enterprise Distribution.

---

## Q

- QDR** Quad data rate. A communication signaling technique wherein data is transmitted at four points in the clock cycle.
- QSFP** Quad small form-factor pluggable. An interconnect system for the I/O ports.

---

## S

- SerDes** Serializer/Deserializer. A pair of functional blocks used in high-speed communications to compensate for limited input/output.

---

## U

- uDAPL** User Direct Access Programming Library.

# Index

---

## A

- acquiring
  - BXOFED software, 18
- adapter
  - features, 3
  - installation, 1
  - LEDs, 6
  - power, 4
  - specifications, 4

## B

- bracket
  - height, 3
  - short bracket removal, 23
  - tall bracket assembly, 25
  - tall bracket installation, 25
- BXOFED
  - acquiring, 18
  - installing, 20

## F

- features, 3
- firmware
  - update for Oracle Solaris, 15

## G

- GUID for node, 8

## H

- hardware and software requirements, 8

## I

- I/O
  - cables, 2
  - panel, 6
- I2C compatible interface, 7

- IB Subnet Manager, 22

## InfiniBand

- devices on Oracle Solaris 10, 16
- interface, 5
- software for Oracle Solaris 10, 11
- installing
  - BXOFED software, 20
- installing the adapter, 1
- interfaces
  - I2C compatible, 7
  - InfiniBand, 5
  - PCI Express, 5
- Internet Protocol Over InfiniBand (IPoIB), 16, 21
- IPoIB
  - Linux, 21
  - Oracle Solaris, 16

## L

- LED, 6, 22
  - assignment, 6
  - location, 6
- Linux
  - distributions, 18
  - IPoIB, 21
  - software installation verification, 21
  - support, 11

## N

- node GUID, 8

## O

- OFED, 17
- Open Fabrics Enterprise Distribution (OFED), 17
- OpenFabrics organization, 18
- Oracle Solaris 10
  - InfiniBand devices, 16

InfiniBand software, 11  
software installation verification, 13

## **P**

PCI Express interface, 5

## **R**

requirements, 8

## **S**

software and hardware requirements, 8

software installation

firmware for Oracle Solaris, 15

Infiniband for Oracle Solaris, 15

Oracle Solaris, 13

verification for Linux, 21

verification for Oracle Solaris, 13

specifications, 4

Subnet Manager, 22

## **U**

uDAPL, 16