

**IB Host Channel Adapter Expansion Card for
IBM eServer BladeCenter User Guide**

Release 3.0.0 Update 2

Copyright © 2004 - 2005 Topspin Communications, Inc. All rights reserved.

The Topspin Switched Computing System, Topspin Host Channel Adapter, Topspin Element Manager, and collateral software programs and documentation are subject to and made available only pursuant to the license agreement signed by you and Topspin, Communications, Inc. (or if no signed license agreement exists, the license agreement included with your original media) and may only be used in accordance with the terms of that agreement. Making copies, modifications, or compilation works of the software except as specifically allowed in that agreement is prohibited by law and constitutes a punishable violation of the law. This documentation may not be copied, modified or reduced to any electronic or machine-readable form without Topspin Communication, Inc.'s prior written consent.

As defined in FAR section 2.101, DFAR section 252.227-7014(a)(1) and DFAR section 252.227-7014(a)(5) or otherwise, the collateral software programs provided to you are “commercial items,” “commercial computer software” and/or “commercial computer software documentation.” Consistent with DFAR section 227.7202, FAR section 12.212 and other sections, any use, modification, reproduction, release, performance, display, disclosure or distribution thereof by or for the U.S. Government shall be governed solely by the terms of the license agreement and shall be prohibited except to the extent expressly permitted by the terms of that agreement. Any technical data provided that is not covered by the above provisions shall be deemed “technical data-commercial items” pursuant to DFAR section 227.7015(a) or equivalent. Any use, modification, reproduction, release, performance, display or disclosure of such technical data shall be governed by the terms of DFAR section 227.7015(b) or equivalent.

This documentation may include technical inaccuracies or typographical errors and is subject to correction and other revision without notice. TOPSPIN PROVIDES THIS DOCUMENTATION “AS IS” AND WITHOUT ANY WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states or jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

©Copyright 2004 - 2005, Topspin Communications, Inc. All rights reserved. Topspin is a registered trademark and the Topspin logo, TopspinOS, and Topspin Switched Computing System are trademarks of Topspin Communications, Inc. Other company, product, or service names are the property of their respective owners.

Document Version: 3.0.0 Update 2

Part Number: 10-00091-02-A0

May, 2005

Printed in the United States of America.

Important Notice

License information for IB Host Channel Adapter Expansion Card for IBM eServer BladeCenter

IMPORTANT NOTICE: This hardware product was manufactured for IBM by Topspin Communications and contains software and/or firmware provided to you by Topspin Communications under provided license agreements. By using the enclosed hardware and/or software products, you are agreeing to the terms of these license agreements, and agree that the terms of such agreement are extended to you solely by Topspin Communications.

First Edition (November 2004)

IBM®, Eserver, BladeCenter, and the e-business logo are trademarks of the IBM Corporation in the United States, other countries, or both.

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Statement of Limited Warranty Preamble

Thank you for your purchase. IBM provides hardware warranty service for this product equivalent to that described in the IBM Statement of Limited Warranty (“SOLW”), a copy of which is included (“[IBM Statement of Limited Warranty Z125-4753-08 04/2004](#)” on page xvii), notwithstanding the section entitled, “What this Warranty Does not Cover” of the SOLW. All references to “Machine” in the SOLW apply to this product (except for its firmware or accompanying software) for purposes of the warranty. References to “IBM’s Official Published Specifications” mean the specifications for the product designated by IBM. Any firmware or software included in or which accompanies this product is subject to the terms and conditions of the manufacturer’s end user license agreement.”

Contents

Important Notice iii

Statement of Limited Warranty Preamble iii

Regulatory Notices xi

FCC Statement xi

Safety Information xi

 Electrical Cautions xi

 General Cautions xii

IBM Support Contact Information xii

 Before You Call xii

Topspin Contact Information xiv

IBM Statement of Limited Warranty Z125-4753-08 04/2004 xvii

Part 1 - General Terms xvii

 What this Warranty Covers xvii

 What this Warranty Does not Cover xviii

 How to Obtain Warranty Service xviii

 What IBM Will Do to Correct Problems xviii

 Exchange of a Machine or Part xix

 Your Additional Responsibilities xix

 Limitation of Liability xix

 Governing Law xx

 Jurisdiction xx

Part 2 - Country-unique Terms xx

 AMERICAS xx

 NORTH AMERICA xxi

ASIA PACIFIC	xxii
EUROPE, MIDDLE EAST, AFRICA (EMEA)	xxiv
EUROPEAN UNION (EU)	xxvi
CONSUMERS	xxvi
THE FOLLOWING TERMS APPLY TO THE COUNTRY SPECIFIED:	xxvii
Part 3 - Warranty Information	xxx
Warranty Period	xxx
Types of Warranty Service	xxxii
EU Country Telephone List	xxxii

Notices xxxiii

Edition notice	xxxiv
Trademarks	xxxiv
Important notes	xxxv
Product recycling and disposal	xxxv
Battery return program	xxxv
Electronic emission notices	xxxvi
Federal Communications Commission (FCC) statement	xxxvi
Industry Canada Class A emission compliance statement	xxxvi
Australia and New Zealand Class A statement	xxxvii
United Kingdom telecommunications safety requirement	xxxvii
European Union EMC Directive conformance statement	xxxvii
Taiwanese Class A warning statement	xxxvii
Chinese Class A warning statement	xxxvii
Japanese Voluntary Control Council for Interference (VCCI) statement	xxxviii

Safety xxxix

1: About the Topspin BladeCenter HCA 1

Introduction	1
Features	2
Supported Blade Servers	2
HCA Expansion Cards and Server Switch Modules	2
Single-Switch Topology	2
Dual-Switch Topology	2
Supported Protocols	3
IPoIB	3
SDP	3
uDAPL	3
SRP	3
MPI	4

2: Installing the HCA Expansion Card 5

Install the HCA Expansion Card Hardware.....	5
Install Multiple HCAs in the HS40 8839 Blade.....	6
Install Two HCAs in an HS20 8843 Blade with a Blade Server Extension.....	14
Prepare to Install Drivers.....	17
Install Linux Drivers.....	17
Install Windows 2000 Drivers.....	17
Install Non-WHQL Drivers on a Windows 2000 Server (both BoIB and non-BoIB).....	18
Verify the Installation.....	20
Install Windows 2003 Drivers.....	21
Install WHQL Drivers on a Windows 2000 or 2003 Server (both BoIB and non-BoIB).....	25
Upgrade Firmware.....	25
Upgrade Firmware with Linux.....	25
Upgrade Firmware with Windows.....	26
Uninstall Windows Host Drivers (2000 or 2003).....	27
Configure High-Availability Ports on Linux Hosts.....	27
Merge Physical Ports.....	27
Unmerge Physical Ports.....	28
Configure High-Availability Ports on Windows Hosts.....	28
Undo HA Ports.....	34

3: Configuring Protocols..... 35

Introduction.....	35
Configure IPoIB.....	35
Configure IPoIB on Linux.....	36
Configure IPoIB on Windows.....	41
Configure SRP.....	45
View SRP Devices on Linux.....	46
View SRP Devices on Windows.....	46
Configure MPI (Linux Only).....	47
Configure SSH.....	47
Edit PATH Variable.....	49
Perform Bandwidth Test.....	49
Perform Latency Test.....	50
Configure SDP (Linux Only).....	50
Prepare.....	50
Convert TCP to SDP in Applications.....	51

4: HCA Utilities..... 53

Introduction.....	53
Run a Self Test.....	54
Run the vstat Utility to View HCA Attributes.....	55
View Host GUID with vstat Utility on Linux.....	55
Run the Basic vstat Utility.....	56
vstat on Windows.....	57
Enable or Disable an Interface.....	57

Linux	58
Windows	58
Run the tvflash -o Utility	58
Configure Boot Type	59
View and Control an HCA Port on Windows.....	60
View the GID of an HCA Port.....	62
View the Pkeys to Which a Port Belongs	63
View InfiniBand (IB) Statistics	64

5: Booting a Host from a Fibre Channel SAN 69

Overview.....	69
Hardware Setup.....	70
Install Hardware.....	70
Check HCA Firmware Version.....	71
Check Version on Linux	71
Check Version on Windows	71
BoIB Prerequisites for Windows	71
Configure a Fibre Channel Connection	72
Use Element Manager and a Host with a Remote Boot HCA	72
Use the CLI and a Host with a Remote Boot HCA	73
Install a Linux Image onto Fibre Channel Storage.....	79
Install and Boot Red Hat 9, Red Hat Enterprise Linux 3.0 or Red Hat Enterprise Linux 4.0	79
Install and Boot Red Hat Enterprise Linux 2.1	86
Install and Boot SUSE Linux Enterprise Server 8 Service Pack 3	88
Install and Boot SUSE Linux Enterprise Server 9 Service Pack 1	94
Boot an Existing Linux Image	101
Boot the Installation Host	101
Boot a New, Diskless Host with a New HCA	101
BoIB Prerequisites for Windows 2000	106
Install an Image onto Fibre Channel Storage: Overview.....	107
Configuration Sample.....	107

6: Booting a Host from PXE 111

About Booting Over PXE.....	111
Sample PXE Hardware Setup.....	111
Overview of the PXE Process.....	111
Configure Boot Over PXE.....	112
Verify Prerequisites	112
Perform Hardware Setup	112
Verify PXE-Enabled Firmware for Booting Server.....	113
Configure the Booting Host.....	113
Configure the DHCP Server	114
Configure the TFTP Server.....	115
Configure Boot Loader	116
Configure Topspin Ethernet Gateway	117

7: Specifications and Compliance 119

Physical and Environmental Specifications.....	119
Electromagnetic Compatibility	120

Regulatory Notices

FCC Statement

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 not installed and used in accordance with the instruction manual, 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.

Safety Information

- [“Electrical Cautions” on page xi](#)
- [“General Cautions” on page xii](#)

Electrical Cautions



CAUTION: Always ground yourself before touching any HCA module component to avoid damage from electrostatic discharge (ESD).

If one or more of the following situations occurs, do not install the HCA Module in a chassis and contact your Topspin support representative:

- The HCA Module has been exposed to water.
- The HCA Module has been dropped or damaged in any way.

- The HCA Module does not operate correctly after following the installation instructions.



CAUTION: You must ensure the operating environment meets the HCA Module's electrical and environmental operating requirements.



CAUTION: Use the Topspin HCA Module only with approved chassis.

General Cautions

In general, do not spill food or liquids on your HCA Module components.

IBM Support Contact Information

For technical support, refer to the following Web site:

<http://www.ibm.com/support>

To contact technical support, send email to support@ibm.com or refer to one of the telephone numbers below.

To obtain warranty service contact IBM or your IBM reseller. In Canada or the United States, call 1-800-IBM-SERV (426-7378). In the EU countries, see the telephone numbers in table [Table iii-1](#).

Table iii-1: EU Country Telephone List

Austria -- 43-1-24592-5901	Italy -- 39-02-482-9202
Belgium -- 02-718-4339	Luxembourg -- 352-360385-1
Denmark -- 4520-8200	Netherlands -- 020-514-5770
Finland -- 358-9-4591	Portugal -- 351-21-7915-147
France -- 0238-557-450	Spain -- 34-91-662-4916
Germany -- 07032-15-4920	Sweden -- 46-8-477-4420
Greece -- 30-210-688-1220	United Kingdom -- 01475-555-055
Ireland -- 353-1-815-4000	

Phone numbers are subject to change without notice.



NOTE: IBM and Topspin are both available for more information. For support, contact IBM before Topspin.

Before You Call

- (Linux only; leave blank for Windows) Enter the **rpm -qa | grep topspin** command at your host to view your RPMs, then provide this output.

- (Linux only; leave blank for Windows) Run an HCA self test.

Example

```
[root@host root]# /usr/local/topspin/sbin/hca_self_test

---- Performing InfiniBand HCA Self Test ----
Number of HCAs Detected ..... 1
PCI Device Check ..... PASS
Host Driver Version ..... rhel3-2.4.21-20.ELsmp-3.0.0-179 rhel3-2
.4.21-4.ELsmp-3.0.0-179
Host Driver RPM Check ..... PASS
HCA Type of HCA #0 ..... Cougar
HCA Firmware on HCA #0 ..... v3.02.0000 build 3.1.0.50 HCA.Cougar.A1
HCA Firmware Check on HCA #0 ..... PASS
Host Driver Initialization ..... PASS
Number of HCA Ports Active ..... 1
Port State of Port #0 on HCA #0 ..... UP
Port State of Port #1 on HCA #0 ..... DOWN
Error Counter Check ..... PASS
Kernel Syslog Check ..... FAIL
    REASON: Kernel syslog reported: Driver messages
    [MOD_LNX_SDP][_tsSdpInetGetOpt][sdp_inet.c:1632]SOCK: GETSOCKOPT unimpleme
nted option <2>.
----- DONE -----
```

- Provide all relevant messages from system and console logs.
- Provide the output from the **show running-status all** command.
- Provide the firmware version of the host. For a Linux host, refer to [step a](#); for a Windows host, refer to [step b](#).
 - a. For Linux hosts, enter the **/usr/local/topspin/sbin/tvflash -i** command to view the firmware version. In the example that follows, the image is **v3.01.0000 build 2.0.0.491**.

Example

```
# /usr/local/topspin/sbin/tvflash -i
HCA #0: Found MT23108, Cougar, revision A1 (firmware autoupgrade)
    Primary image is v3.01.0000 build 2.0.0.491, with label 'HCA.Cougar.A1.Boot'
    Secondary image is valid, unknown source
[root@enclus2 root]#
```

- a. For Windows hosts, perform the following steps:
 1. Log in to your Windows host.
 2. Click **Start > Programs > Topspin IB SDK > Utilities**. The **Utilities** window opens and displays a list of utilities.
 3. Enter **tvflash -i** command to display the firmware version.

- (Linux only) Run the **dmesg** command on your Linux host and provide the output.

Example

```
# dmesg
 0  0  0  1  1  71
0a 000 00 1 0 0 0 0 0 0 00
0b 000 00 1 0 0 0 0 0 0 00
0c 00F 0F 0 0 0 0 0 1 1 79
0d 00F 0F 0 0 0 0 0 1 1 81
0e 00F 0F 0 0 0 0 0 1 1 89
0f 00F 0F 0 0 0 0 0 1 1 91
10 00F 0F 1 1 0 1 0 1 1 99
11 000 00 1 0 0 0 0 0 0 00
12 00F 0F 1 1 0 1 0 1 1 A1
13 00F 0F 1 1 0 1 0 1 1 A9
14 000 00 1 0 0 0 0 0 0 00
15 000 00 1 0 0 0 0 0 0 00
16 000 00 1 0 0 0 0 0 0 00
17 000 00 1 0 0 0 0 0 0 00

//content truncated

scsi : 2 hosts left.
scsi2 : srp
initializing plug timer for queue f69fcel8
Starting timer : 1 1
blk: queue f69fcel8, I/O limit 4095Mb (mask 0xffffffff)
  Vendor: SEAGATE    Model: ST336607FC    Rev: 0004
  Type:   Direct-Access    ANSI SCSI revision: 03
initializing plug timer for queue f69fd018
Starting timer : 1 1
blk: queue f69fd018, I/O limit 4095Mb (mask 0xffffffff)
Attached scsi disk sdb at scsi2, channel 0, id 6, lun 0
SCSI device sdb: 71687372 512-byte hdwr sectors (36704 MB)
  sdb: sdb1 sdb2
srp_host:
target_bindings=22000004cfe7a949.0:22000004cfe7aa7a.1:22000004cfe7aac7.2:2200000
4cfe7aadd.3:22000004cfe7ab01.4:22000004cfe7ab2f.5:22000004cfe7ab35.6:22000004cfe
7ab6c.7:22000004cfe7ab7b.8:22000004cfe7ab89.9:22000004cfe7ab8d.a:22000004cfe7ac1
f.b:22000004cfe7ac30.c:22000004cfe7ac69.d
divert: allocating divert_blk for ib0
divert: allocating divert_blk for ib1
nfs: server 10.5 not responding, still trying
nfs: server 10.5 OK
```

Topspin Contact Information

Topspin contact information appears in [Table iii-2](#).

Table iii-2: Contact Information

Address	Topspin Communications, Inc. 515 Ellis St. Mountain View, CA 94043
---------	--

Table iii-2: Contact Information (Continued)

Telephone, Corporate Headquarters	(866) TOPSPIN or (650) 316-3300
Fax	(650) 316-3269
Telephone, Technical Support	(800) 499-1473
Email, Technical Support	support@ibm.com (contact IBM first) support@topspin.com
Web site, Support	http://www.ibm.com/support (contact IBM first) http://support.topspin.com

IBM Statement of Limited Warranty

Z125-4753-08 04/2004

Part 1 - General Terms

Part 1 - General Terms

This Statement of Limited Warranty includes Part 1 - General Terms, Part 2 - Country-unique Terms, and Part 3 - Warranty Information. The terms of Part 2 replace or modify those of Part 1. The warranties provided by IBM in this Statement of Limited Warranty apply only to Machines you purchase for your use, and not for resale. The term "Machine" means an IBM machine, its features, conversions, upgrades, elements, or accessories, or any combination of them. The term "Machine" does not include any software programs, whether pre-loaded with the Machine, installed subsequently or otherwise. Nothing in this Statement of Limited Warranty affects any statutory rights of consumers that cannot be waived or limited by contract.

What this Warranty Covers

IBM warrants that each Machine 1) is free from defects in materials and workmanship and 2) conforms to IBM's Official Published Specifications ("Specifications") which are available on request. The warranty period for the Machine starts on the original Date of Installation and is specified in Part 3 - Warranty Information. The date on your invoice or sales receipt is the Date of Installation unless IBM or your reseller informs you otherwise. Many features, conversions, or upgrades involve the removal of parts and their return to IBM. A part that replaces a removed part will assume the warranty service status of the removed part. Unless IBM specifies otherwise, these warranties apply only in the country or region in which you purchased the Machine.

THESE WARRANTIES ARE YOUR EXCLUSIVE WARRANTIES AND REPLACE ALL OTHER WARRANTIES OR CONDITIONS, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. SOME STATES OR

JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF EXPRESS OR IMPLIED WARRANTIES, SO THE ABOVE EXCLUSION MAY NOT APPLY TO YOU. IN THAT EVENT, SUCH WARRANTIES ARE LIMITED IN DURATION TO THE WARRANTY PERIOD. NO WARRANTIES APPLY AFTER THAT PERIOD. SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

What this Warranty Does not Cover

This warranty does not cover the following:

- any software programs, whether pre-loaded or shipped with the Machine, or installed subsequently;
- failure resulting from misuse (including but not limited to use of any Machine capacity or capability, other than that authorized by IBM in writing), accident, modification, unsuitable physical or operating environment, or improper maintenance by you;
- failure caused by a product for which IBM is not responsible; and
- any non-IBM products, including those that IBM may procure and provide with or integrate into an IBM Machine at your request.

The warranty is voided by removal or alteration of identification labels on the Machine or its parts.

IBM does not warrant uninterrupted or error-free operation of a Machine.

Any technical or other support provided for a Machine under warranty, such as assistance with “how-to” questions and those regarding Machine set-up and installation, is provided **WITHOUT WARRANTIES OF ANY KIND**.

How to Obtain Warranty Service

If the Machine does not function as warranted during the warranty period, contact IBM or your reseller to obtain warranty service. If you do not register the Machine with IBM, you may be required to present proof of purchase as evidence of your entitlement to warranty service.

What IBM Will Do to Correct Problems

When you contact IBM for service, you must follow the problem determination and resolution procedures that IBM specifies. An initial diagnosis of your problem can be made either by a technician over the telephone or electronically by access to an IBM Web site.

The type of warranty service applicable to your Machine is specified in Part 3 - Warranty Information.

You are responsible for downloading and installing designated Machine Code (microcode, basic input/output system code (called “BIOS”), utility programs, device drivers, and diagnostics delivered with an IBM Machine) and other software updates from an IBM Internet Web site or from other electronic media, and following the instructions that IBM provides.

If your problem can be resolved with a Customer Replaceable Unit (“CRU”) (e.g., keyboard, mouse, speaker, memory, hard disk drive), IBM will ship the CRU to you for you to install.

If the Machine does not function as warranted during the warranty period and your problem cannot be resolved over the telephone or electronically, through your application of Machine Code or software updates, or with a CRU, IBM or your reseller, if approved by IBM to provide warranty service, will either, at its discretion, 1) repair it to make it function as warranted, or 2) replace it with one that is at least functionally equivalent. If IBM is unable to do either, you may return the Machine to your place of purchase and your money will be refunded.

IBM or your reseller will also manage and install selected engineering changes that apply to the Machine.

Exchange of a Machine or Part

When the warranty service involves the exchange of a Machine or part, the item IBM or your reseller replaces becomes its property and the replacement becomes yours. You represent that all removed items are genuine and unaltered. The replacement may not be new, but will be in good working order and at least functionally equivalent to the item replaced. The replacement assumes the warranty service status of the replaced item.

Your Additional Responsibilities

Before IBM or your reseller exchanges a Machine or part, you agree to remove all features, parts, options, alterations, and attachments not under warranty service.

You also agree to:

1. ensure that the Machine is free of any legal obligations or restrictions that prevent its exchange;
2. obtain authorization from the owner to have IBM or your reseller service a Machine that you do not own; and
3. where applicable, before service is provided:
 - a. follow the service request procedures that IBM or your reseller provides;
 - b. backup or secure all programs, data, and funds contained in the Machine;
 - c. provide IBM or your reseller with sufficient, free, and safe access to your facilities to permit IBM to fulfill its obligations; and
 - d. inform IBM or your reseller of changes in the Machine's location.
4. (a) ensure all information about identified or identifiable individuals (Personal Data) is deleted from the Machine (to the extent technically possible), (b) allow IBM, your reseller or an IBM supplier to process on your behalf any remaining Personal Data as IBM or your reseller considers necessary to fulfill its obligations under this Statement of Limited Warranty (which may include shipping the Machine for such processing to other IBM service locations around the world), and (c) ensure that such processing complies with any laws applicable to such Personal Data.

Limitation of Liability

IBM is responsible for loss of, or damage to, your Machine only while it is 1) in IBM's possession or 2) in transit in those cases where IBM is responsible for the transportation charges.

Neither IBM nor your reseller are responsible for any of your confidential, proprietary or personal information contained in a Machine which you return to IBM for any reason. You should remove all such information from the Machine prior to its return.

Circumstances may arise where, because of a default on IBM's part or other liability, you are entitled to recover damages from IBM. In each such instance, regardless of the basis on which you are entitled to claim damages from IBM (including fundamental breach, negligence, misrepresentation, or other contract or tort claim), except for any liability that cannot be waived or limited by applicable laws, IBM is liable for no more than

1. damages for bodily injury (including death) and damage to real property and tangible personal property; and
2. the amount of any other actual direct damages, up to the charges (if recurring, 12 months' charges apply) for the Machine that is subject of the claim. For purposes of this item, the term "Machine" includes Machine Code and Licensed Internal Code ("LIC").

This limit also applies to IBM's suppliers and your reseller. It is the maximum for which IBM, its suppliers, and your reseller are collectively responsible.

UNDER NO CIRCUMSTANCES IS IBM, ITS SUPPLIERS OR RESELLERS LIABLE FOR ANY OF THE FOLLOWING EVEN IF INFORMED OF THEIR POSSIBILITY: 1) THIRD

PARTY CLAIMS AGAINST YOU FOR DAMAGES (OTHER THAN THOSE UNDER THE FIRST ITEM LISTED ABOVE); 2) LOSS OF, OR DAMAGE TO, DATA; 3) SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR 4) LOST PROFITS, BUSINESS REVENUE, GOODWILL OR ANTICIPATED SAVINGS. SOME STATES OR JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION OR EXCLUSION MAY NOT APPLY TO YOU. SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

Governing Law

Both you and IBM consent to the application of the laws of the country in which you acquired the Machine to govern, interpret, and enforce all of your and IBM's rights, duties, and obligations arising from, or relating in any manner to, the subject matter of this Statement of Limited Warranty, without regard to conflict of law principles.

THESE WARRANTIES GIVE YOU SPECIFIC LEGAL RIGHTS AND YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM STATE TO STATE OR JURISDICTION TO JURISDICTION.

Jurisdiction

All of our rights, duties, and obligations are subject to the courts of the country in which you acquired the Machine.

Part 2 - Country-unique Terms

AMERICAS

ARGENTINA

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Ordinary Commercial Court of the city of Buenos Aires.

BOLIVIA

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the courts of the city of La Paz.

BRAZIL

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the court of Rio de Janeiro, RJ.

CHILE

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Civil Courts of Justice of Santiago.

COLOMBIA

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Judges of the Republic of Colombia.

EQUADOR

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Judges of Quito.

MEXICO

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Federal Courts of Mexico City, Federal District.

PARAGUAY

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the courts of the city of Asuncion.

PERU

Limitation of Liability: The following is added at the end of this section:

In accordance with Article 1328 of the Peruvian Civil Code the limitations and exclusions specified in this section will not apply to damages caused by IBM's willful misconduct ("dolo") or gross negligence ("culpa inexcusable").

URUGUAY

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the City of Montevideo Court's Jurisdiction.

VENEZUELA

Jurisdiction: *The following is added after the first sentence:*

Any litigation arising from this Statement of Limited Warranty will be settled exclusively by the Courts of the Metropolitan Area Of the City of Caracas.

NORTH AMERICA

How to Obtain Warranty Service: *The following is added to this Section:*

To obtain warranty service from IBM in Canada or the United States, call 1-800-IBM-SERV (426-7378).

CANADA

Limitation of Liability: *The following replaces item 1 of this section:*

1. damages for bodily injury (including death) or physical harm to real property and tangible personal property caused by IBM's negligence; and

Governing Law: *The following replaces "laws of the country in which you acquired the Machine" in the first sentence:*

laws in the Province of Ontario.

UNITED STATES

Governing Law: The following replaces "laws of the country in which you acquired the Machine" in the first sentence:

laws of the State of New York.

ASIA PACIFIC

AUSTRALIA

What this Warranty Covers: *The following paragraph is added to this section:*

The warranties specified in this Section are in addition to any rights you may have under the Trade Practices Act 1974 or other similar legislation and are only limited to the extent permitted by the applicable legislation.

Limitation of Liability: *The following is added to this section:*

Where IBM is in breach of a condition or warranty implied by the Trade Practices Act 1974 or other similar legislation, IBM's liability is limited to the repair or replacement of the goods or the supply of equivalent goods. Where that condition or warranty relates to right to sell, quiet possession or clear title, or the goods are of a kind ordinarily acquired for personal, domestic or household use or consumption, then none of the limitations in this paragraph apply.

Governing Law: The following replaces "laws of the country in which you acquired the Machine" in the first sentence:

laws of the State or Territory.

CAMBODIA AND LAOS

Governing Law: *The following replaces "laws of the country in which you acquired the Machine" in the first sentence:*

laws of the State of New York, United States of America.

CAMBODIA, INDONESIA, AND LAOS

Arbitration: *The following is added under this heading:*

Disputes arising out of or in connection with this Statement of Limited Warranty shall be finally settled by arbitration which shall be held in Singapore in accordance with the Arbitration Rules of Singapore International Arbitration Center ("SIAC Rules") then in effect. The arbitration award shall be final and binding for the parties without appeal and shall be in writing and set forth the findings of fact and the conclusions of law.

The number of arbitrators shall be three, with each side to the dispute being entitled to appoint one arbitrator. The two arbitrators appointed by the parties shall appoint a third arbitrator who shall act as chairman of the proceedings. Vacancies in the post of chairman shall be filled by the president of the SIAC. Other vacancies shall be filled by the respective nominating party. Proceedings shall continue from the stage they were at when the vacancy occurred.

If one of the parties refuses or otherwise fails to appoint an arbitrator within 30 days of the date the other party appoints its, the first appointed arbitrator shall be the sole arbitrator, provided that the arbitrator was validly and properly appointed.

All proceedings shall be conducted, including all documents presented in such proceedings, in the English language. The English language version of this Statement of Limited Warranty prevails over any other language version.

HONG KONG S.A.R. OF CHINA AND MACAU S.A.R. OF CHINA

Governing Law: *The following replaces “laws of the country in which you acquired the Machine” in the first sentence:*

laws of Hong Kong Special Administrative Region of China.

INDIA

Limitation of Liability: *The following replaces items 1 and 2 of this section:*

1. liability for bodily injury (including death) or damage to real property and tangible personal property will be limited to that caused by IBM's negligence; and
2. as to any other actual damage arising in any situation involving nonperformance by IBM pursuant to, or in any way related to the subject of this Statement of Limited Warranty, the charge paid by you for the individual Machine that is the subject of the claim. For purposes of this item, the term “Machine” includes Machine Code and Licensed Internal Code (“LIC”).

Arbitration: *The following is added under this heading:*

Disputes arising out of or in connection with this Statement of Limited Warranty shall be finally settled by arbitration which shall be held in Bangalore, India in accordance with the laws of India then in effect. The arbitration award shall be final and binding for the parties without appeal and shall be in writing and set forth the findings of fact and the conclusions of law.

The number of arbitrators shall be three, with each side to the dispute being entitled to appoint one arbitrator. The two arbitrators appointed by the parties shall appoint a third arbitrator who shall act as chairman of the proceedings. Vacancies in the post of chairman shall be filled by the president of the Bar Council of India. Other vacancies shall be filled by the respective nominating party. Proceedings shall continue from the stage they were at when the vacancy occurred.

If one of the parties refuses or otherwise fails to appoint an arbitrator within 30 days of the date the other party appoints its, the first appointed arbitrator shall be the sole arbitrator, provided that the arbitrator was validly and properly appointed.

All proceedings shall be conducted, including all documents presented in such proceedings, in the English language. The English language version of this Statement of Limited Warranty prevails over any other language version.

JAPAN

Governing Law: *The following sentence is added to this section:*

Any doubts concerning this Statement of Limited Warranty will be initially resolved between us in good faith and in accordance with the principle of mutual trust.

MALAYSIA

Limitation of Liability: *The word “SPECIAL” in item 3 of the fifth paragraph is deleted.*

NEW ZEALAND

What this Warranty Covers: *The following paragraph is added to this section:*

The warranties specified in this section are in addition to any rights you may have under the Consumer Guarantees Act 1993 or other legislation which cannot be excluded or limited. The Consumer

Guarantees Act 1993 will not apply in respect of any goods which IBM provides, if you require the goods for the purposes of a business as defined in that Act.

Limitation of Liability: *The following is added to this section:*

Where Machines are not acquired for the purposes of a business as defined in the Consumer Guarantees Act 1993, the limitations in this section are subject to the limitations in that Act.

PEOPLE'S REPUBLIC OF CHINA (PRC)

Governing Law: *The following replaces “laws of the country in which you acquired the Machine” in the first sentence:*

laws of the State of New York, United States of America (except when local law requires otherwise).

PHILIPPINES

Limitation of Liability: *Item 3 in the fifth paragraph is replaced by the following:*

SPECIAL (INCLUDING NOMINAL AND EXEMPLARY DAMAGES), MORAL, INCIDENTAL, OR INDIRECT DAMAGES FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR

Arbitration: *The following is added under this heading:*

Disputes arising out of or in connection with this Statement of Limited Warranty shall be finally settled by arbitration which shall be held in Metro Manila, Philippines in accordance with the laws of the Philippines then in effect. The arbitration award shall be final and binding for the parties without appeal and shall be in writing and set forth the findings of fact and the conclusions of law.

The number of arbitrators shall be three, with each side to the dispute being entitled to appoint one arbitrator. The two arbitrators appointed by the parties shall appoint a third arbitrator who shall act as chairman of the proceedings. Vacancies in the post of chairman shall be filled by the president of the Philippine Dispute Resolution Center, Inc. Other vacancies shall be filled by the respective nominating party. Proceedings shall continue from the stage they were at when the vacancy occurred.

If one of the parties refuses or otherwise fails to appoint an arbitrator within 30 days of the date the other party appoints its, the first appointed arbitrator shall be the sole arbitrator, provided that the arbitrator was validly and properly appointed.

All proceedings shall be conducted, including all documents presented in such proceedings, in the English language. The English language version of this Statement of Limited Warranty prevails over any other language version.

SINGAPORE

Limitation of Liability: *The words “SPECIAL” and “ECONOMIC” in item 3 in the fifth paragraph are deleted.*

EUROPE, MIDDLE EAST, AFRICA (EMEA)

THE FOLLOWING TERMS APPLY TO ALL EMEA COUNTRIES:

The terms of this Statement of Limited Warranty apply to Machines purchased from IBM or an IBM reseller.

How to Obtain Warranty Service:

*Add the following paragraph in **Western Europe** (Andorra, Austria, Belgium, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Netherlands, Norway, Poland, Portugal, San Marino, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Kingdom, Vatican State, and any country subsequently added to the European Union, as from the date of accession):*

The warranty for Machines acquired in Western Europe shall be valid and applicable in all Western Europe countries provided the Machines have been announced and made available in such countries.

If you purchase a Machine in one of the Western European countries, as defined above, you may obtain warranty service for that Machine in any of those countries from either (1) an IBM reseller approved to perform warranty service or (2) from IBM, provided the Machine has been announced and made available by IBM in the country in which you wish to obtain service.

If you purchased a Personal Computer Machine in Albania, Armenia, Belarus, Bosnia and Herzegovina, Bulgaria, Croatia, Czech Republic, Georgia, Hungary, Kazakhstan, Kyrgyzstan, Federal Republic of Yugoslavia, Former Yugoslav Republic of Macedonia (FYROM), Moldova, Poland, Romania, Russia, Slovak Republic, Slovenia, or Ukraine, you may obtain warranty service for that Machine in any of those countries from either (1) an IBM reseller approved to perform warranty service or (2) from IBM.

If you purchase a Machine in a Middle Eastern or African country, you may obtain warranty service for that Machine from the IBM entity within the country of purchase, if that IBM entity provides warranty service in that country, or from an IBM reseller, approved by IBM to perform warranty service on that Machine in that country. Warranty service in Africa is available within 50 kilometers of an IBM approved service provider. You are responsible for transportation costs for Machines located outside 50 kilometers of an IBM approved service provider.

Governing Law:

The phrase “the laws of the country in which you acquired the Machine” is replaced by:

1) “the laws of Austria” in **Albania, Armenia, Azerbaijan, Belarus, Bosnia-Herzegovina, Bulgaria, Croatia, Georgia, Hungary, Kazakhstan, Kyrgyzstan, FYR Macedonia, Moldova, Poland, Romania, Russia, Slovakia, Slovenia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan, and FR Yugoslavia**; 2) “the laws of France” in **Algeria, Benin, Burkina Faso, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Congo Republic, Djibouti, Democratic Republic of Congo, Equatorial Guinea, French Guiana, French Polynesia, Gabon, Gambia, Guinea, Guinea-Bissau, Ivory Coast, Lebanon, Madagascar, Mali, Mauritania, Mauritius, Mayotte, Morocco, New Caledonia, Niger, Reunion, Senegal, Seychelles, Togo, Tunisia, Vanuatu, and Wallis & Futuna**; 3) “the laws of Finland” in **Estonia, Latvia, and Lithuania**; 4) “the laws of England” in **Angola, Bahrain, Botswana, Burundi, Egypt, Eritrea, Ethiopia, Ghana, Jordan, Kenya, Kuwait, Liberia, Malawi, Malta, Mozambique, Nigeria, Oman, Pakistan, Qatar, Rwanda, Sao Tome, Saudi Arabia, Sierra Leone, Somalia, Tanzania, Uganda, United Arab Emirates, the United Kingdom, West Bank/Gaza, Yemen, Zambia, and Zimbabwe**; and 5) “the laws of South Africa” in **South Africa, Namibia, Lesotho and Swaziland**.

Jurisdiction: *The following is added after the first sentence:*

1) **In Austria** the choice of jurisdiction for all disputes arising out of this Statement of Limited Warranty and relating thereto, including its existence, will be the competent court of law in Vienna, Austria (Inner-City); 2) **in Angola, Bahrain, Botswana, Burundi, Egypt, Eritrea, Ethiopia, Ghana, Jordan, Kenya, Kuwait, Liberia, Malawi, Malta, Mozambique, Nigeria, Oman, Pakistan, Qatar, Rwanda, Sao Tome, Saudi Arabia, Sierra Leone, Somalia, Tanzania, Uganda, United Arab Emirates, West Bank/Gaza, Yemen, Zambia, and Zimbabwe** all disputes arising out of this Statement of Limited Warranty or related to its execution, including summary proceedings, will be submitted to the exclusive jurisdiction of the English courts; 3) **in Belgium and Luxembourg**, all disputes arising out of this Statement of Limited Warranty or related to its interpretation or its execution, the law, and the courts of the capital city, of the country of your registered office and/or commercial site location only are competent; 4) **in France, Algeria, Benin, Burkina Faso, Cameroon, Cape Verde, Central African Republic, Chad, Comoros, Congo Republic, Djibouti, Democratic Republic of Congo, Equatorial Guinea, French Guiana, French Polynesia, Gabon, Gambia, Guinea, Guinea-Bissau, Ivory Coast, Lebanon, Madagascar, Mali, Mauritania, Mauritius, Mayotte, Morocco, New Caledonia, Niger, Reunion, Senegal, Seychelles, Togo, Tunisia, Vanuatu, and Wallis & Futuna** all disputes arising out of this Statement of Limited Warranty or related to its violation or execution, including summary proceedings, will be settled exclusively by the Commercial Court of Paris; 5) **in Russia**, all disputes arising out of or in relation to the interpretation, the violation, the termination, the nullity of the execution of this Statement of Limited Warranty shall be settled by

Arbitration Court of Moscow; 6) **in South Africa, Namibia, Lesotho and Swaziland**, both of us agree to submit all disputes relating to this Statement of Limited Warranty to the jurisdiction of the High Court in Johannesburg; 7) **in Turkey** all disputes arising out of or in connection with this Statement of Limited Warranty shall be resolved by the Istanbul Central (Sultanahmet) Courts and Execution Directorates of Istanbul, the Republic of Turkey; 8) in each of the following specified countries, any legal claim arising out of this Statement of Limited Warranty will be brought before, and settled exclusively by, the competent court of a) Athens for **Greece**, b) Tel Aviv-Jaffa for **Israel**, c) Milan for **Italy**, d) Lisbon for **Portugal**, and e) Madrid for **Spain**; and 9) **in the United Kingdom**, both of us agree to submit all disputes relating to this Statement of Limited Warranty to the jurisdiction of the English courts.

Arbitration: *The following is added under this heading:*

In Albania, Armenia, Azerbaijan, Belarus, Bosnia-Herzegovina, Bulgaria, Croatia, Georgia, Hungary, Kazakhstan, Kyrgyzstan, FYR Macedonia, Moldova, Poland, Romania, Russia, Slovakia, Slovenia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan, and FR Yugoslavia all disputes arising out of this Statement of Limited Warranty or related to its violation, termination or nullity will be finally settled under the Rules of Arbitration and Conciliation of the International Arbitral Center of the Federal Economic Chamber in Vienna (Vienna Rules) by three arbitrators appointed in accordance with these rules. The arbitration will be held in Vienna, Austria, and the official language of the proceedings will be English. The decision of the arbitrators will be final and binding upon both parties. Therefore, pursuant to paragraph 598 (2) of the Austrian Code of Civil Procedure, the parties expressly waive the application of paragraph 595 (1) figure 7 of the Code. IBM may, however, institute proceedings in a competent court in the country of installation.

In Estonia, Latvia and Lithuania all disputes arising in connection with this Statement of Limited Warranty will be finally settled in arbitration that will be held in Helsinki, Finland in accordance with the arbitration laws of Finland then in effect. Each party will appoint one arbitrator. The arbitrators will then jointly appoint the chairman. If arbitrators cannot agree on the chairman, then the Central Chamber of Commerce in Helsinki will appoint the chairman.

EUROPEAN UNION (EU)

THE FOLLOWING TERMS APPLY TO ALL EU COUNTRIES:

The warranty for Machines acquired in EU countries is valid and applicable in all EU countries provided the Machines have been announced and made available in such countries.

How to Obtain Warranty Service: *The following is added to this section:*

To obtain warranty service from IBM in EU countries, see the telephone listing in Part 3 - Warranty Information.

You may contact IBM at the following address:

IBM Warranty & Service Quality Dept.
PO Box 30
Spango Valley
Greenock
Scotland PA16 0AH

CONSUMERS

Consumers have legal rights under applicable national legislation governing the sale of consumer goods. Such rights are not affected by the warranties provided in this Statement of Limited Warranty.

AUSTRIA, DENMARK, FINLAND, GREECE, ITALY, NETHERLANDS, NORWAY, PORTUGAL, SPAIN, SWEDEN AND SWITZERLAND

Limitation of Liability: *The following replaces the terms of this section in its entirety:*

Except as otherwise provided by mandatory law:

1. IBM's liability for any damages and losses that may arise as a consequence of the fulfillment of its obligations under or in connection with this Statement of Limited Warranty or due to any other cause related to this Statement of Limited Warranty is limited to the compensation of only those damages and losses proved and actually arising as an immediate and direct consequence of the non-fulfillment of such obligations (if IBM is at fault) or of such cause, for a maximum amount equal to the charges you paid for the Machine. For purposes of this item, the term "Machine" includes Machine Code and Licensed Internal Code ("LIC").

The above limitation shall not apply to damages for bodily injuries (including death) and damages to real property and tangible personal property for which IBM is legally liable.

2. **UNDER NO CIRCUMSTANCES IS IBM, ITS SUPPLIERS OR RESELLERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY: 1) LOSS OF, OR DAMAGE TO, DATA; 2) INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; 3) LOST PROFITS, EVEN IF THEY ARISE AS AN IMMEDIATE CONSEQUENCE OF THE EVENT THAT GENERATED THE DAMAGES; OR 4) LOSS OF BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.**

FRANCE AND BELGIUM

Limitation of Liability: *The following replaces the terms of this section in its entirety:*

Except as otherwise provided by mandatory law:

1. IBM's liability for any damages and losses that may arise as a consequence of the fulfillment of its obligations under or in connection with this Statement of Limited Warranty is limited to the compensation of only those damages and losses proved and actually arising as an immediate and direct consequence of the non-fulfillment of such obligations (if IBM is at fault), for a maximum amount equal to the charges you paid for the Machine that has caused the damages. For purposes of this item, the term "Machine" includes Machine Code and Licensed Internal Code ("LIC").

The above limitation shall not apply to damages for bodily injuries (including death) and damages to real property and tangible personal property for which IBM is legally liable.

2. **UNDER NO CIRCUMSTANCES IS IBM, ITS SUPPLIERS OR RESELLERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY: 1) LOSS OF, OR DAMAGE TO, DATA; 2) INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; 3) LOST PROFITS, EVEN IF THEY ARISE AS AN IMMEDIATE CONSEQUENCE OF THE EVENT THAT GENERATED THE DAMAGES; OR 4) LOSS OF BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.**

THE FOLLOWING TERMS APPLY TO THE COUNTRY SPECIFIED:

AUSTRIA

The provisions of this Statement of Limited Warranty replace any applicable statutory warranties.

What this Warranty Covers: *The following replaces the first sentence of the first paragraph of this section:*

The warranty for an IBM Machine covers the functionality of the Machine for its normal use and the Machine's conformity to its Specifications.

The following paragraphs are added to this section:

The limitation period for consumers in action for breach of warranty is the statutory period as a minimum. In case IBM or your reseller is unable to repair an IBM Machine, you can alternatively ask for a partial refund as far as justified by the reduced value of the unrepaired Machine or ask for a cancellation of the respective agreement for such Machine and get your money refunded.

The second paragraph does not apply.

What IBM Will Do to Correct Problems: *The following is added to this section:*

During the warranty period, IBM will reimburse you for the transportation charges for the delivery of the failing Machine to IBM.

Limitation of Liability: *The following paragraph is added to this section:*

The limitations and exclusions specified in the Statement of Limited Warranty will not apply to damages caused by IBM with fraud or gross negligence and for express warranty.

The following sentence is added to the end of item 2:

IBM's liability under this item is limited to the violation of essential contractual terms in cases of ordinary negligence.

EGYPT

Limitation of Liability: *The following replaces item 2 in this section:*

as to any other actual direct damages, IBM's liability will be limited to the total amount you paid for the Machine that is the subject of the claim. For purposes of this item, the term "Machine" includes Machine Code and Licensed Internal Code ("LIC").

Applicability of suppliers and resellers (unchanged).

FRANCE

Limitation of Liability: *The following replaces the second sentence of the first paragraph of this section:*

In such instances, regardless of the basis on which you are entitled to claim damages from IBM, IBM is liable for no more than: (items 1 and 2 unchanged).

GERMANY

What this Warranty Covers: *The following replaces the first sentence of the first paragraph of this section:*

The warranty for an IBM Machine covers the functionality of the Machine for its normal use and the Machine's conformity to its Specifications.

The following paragraphs are added to this section:

The minimum warranty period for Machines is twelve months. In case IBM or your reseller is unable to repair an IBM Machine, you can alternatively ask for a partial refund as far as justified by the reduced value of the unrepaired Machine or ask for a cancellation of the respective agreement for such Machine and get your money refunded.

The second paragraph does not apply.

What IBM Will Do to Correct Problems: *The following is added to this section:*

During the warranty period, transportation for delivery of the failing Machine to IBM will be at IBM's expense.

Limitation of Liability: *The following paragraph is added to this section:*

The limitations and exclusions specified in the Statement of Limited Warranty will not apply to damages caused by IBM with fraud or gross negligence and for express warranty.

The following sentence is added to the end of item 2:

IBM's liability under this item is limited to the violation of essential contractual terms in cases of ordinary negligence.

HUNGARY

Limitation of Liability: *The following is added at the end of this section:*

The limitation and exclusion specified herein shall not apply to liability for a breach of contract damaging life, physical well-being, or health that has been caused intentionally, by gross negligence, or by a criminal act.

The parties accept the limitations of liability as valid provisions and state that the Section 314.(2) of the Hungarian Civil Code applies as the acquisition price as well as other advantages arising out of the present Statement of Limited Warranty balance this limitation of liability.

IRELAND

What this Warranty Covers: *The following is added to this section:*

Except as expressly provided in these terms and conditions, all statutory conditions, including all warranties implied, but without prejudice to the generality of the foregoing all warranties implied by the Sale of Goods Act 1893 or the Sale of Goods and Supply of Services Act 1980 are hereby excluded.

Limitation of Liability: *The following replaces the terms of this section in its entirety:*

For the purposes of this section, a "Default" means any act, statement, omission, or negligence on the part of IBM in connection with, or in relation to, the subject matter of this Statement of Limited Warranty in respect of which IBM is legally liable to you, whether in contract or tort. A number of Defaults which together result in, or contribute to, substantially the same loss or damage will be treated as one Default occurring on the date of occurrence of the last such Default.

Circumstances may arise where, because of a Default, you are entitled to recover damages from IBM.

This section sets out the extent of IBM's liability and your sole remedy.

1. IBM will accept unlimited liability for death or personal injury caused by the negligence of IBM.
2. Subject always to the Items for Which IBM is Not Liable below, IBM will accept unlimited liability for physical damage to your tangible property resulting from the negligence of IBM.
3. Except as provided in items 1 and 2 above, IBM's entire liability for actual damages for any one Default will not in any event exceed the greater of 1) EUR 125,000, or 2) 125% of the amount you paid for the Machine directly relating to the Default.

Items for Which IBM is Not Liable

Save with respect to any liability referred to in item 1 above, under no circumstances is IBM, its suppliers or resellers liable for any of the following, even if IBM or they were informed of the possibility of such losses:

1. loss of, or damage to, data;
2. special, indirect, or consequential loss; or
3. loss of profits, business, revenue, goodwill, or anticipated savings.

SLOVAKIA

Limitation of Liability: *The following is added to the end of the last paragraph:*

The limitations apply to the extent they are not prohibited under §§ 373-386 of the Slovak Commercial Code.

SOUTH AFRICA, NAMIBIA, BOTSWANA, LESOTHO AND SWAZILAND

Limitation of Liability: *The following is added to this section:*

IBM's entire liability to you for actual damages arising in all situations involving nonperformance by IBM in respect of the subject matter of this Statement of Warranty will be limited to the charge paid by you for the individual Machine that is the subject of your claim from IBM.

UNITED KINGDOM

Limitation of Liability: *The following replaces the terms of this section in its entirety:*

For the purposes of this section, a "Default" means any act, statement, omission, or negligence on the part of IBM in connection with, or in relation to, the subject matter of this Statement of Limited Warranty in respect of which IBM is legally liable to you, whether in contract or tort. A number of Defaults which together result in, or contribute to, substantially the same loss or damage will be treated as one Default.

Circumstances may arise where, because of a Default, you are entitled to recover damages from IBM.

This section sets out the extent of IBM's liability and your sole remedy.

1. IBM will accept unlimited liability for:
 - a. death or personal injury caused by the negligence of IBM; and
 - b. any breach of its obligations implied by Section 12 of the Sale of Goods Act 1979 or Section 2 of the Supply of Goods and Services Act 1982, or any statutory modification or re-enactment of either such Section.
2. IBM will accept unlimited liability, subject always to the Items for Which IBM is Not Liable below, for physical damage to your tangible property resulting from the negligence of IBM.
3. IBM's entire liability for actual damages for any one Default will not in any event, except as provided in items 1 and 2 above, exceed the greater of 1) Pounds Sterling 75,000, or 2) 125% of the total purchase price payable or the charges for the Machine directly relating to the Default.

These limits also apply to IBM's suppliers and resellers. They state the maximum for which IBM and such suppliers and resellers are collectively responsible.

Items for Which IBM is Not Liable

Save with respect to any liability referred to in item 1 above, under no circumstances is IBM or any of its suppliers or resellers liable for any of the following, even if IBM or they were informed of the possibility of such losses:

1. loss of, or damage to, data;
2. special, indirect, or consequential loss; or
3. loss of profits, business, revenue, goodwill, or anticipated savings.

Part 3 - Warranty Information

This Part 3 provides information regarding the warranty applicable to your Machine, including the warranty period and type of warranty service IBM provides.

Warranty Period

The warranty period may vary by country or region and is specified in the table below.



NOTE: “Region” means either Hong Kong or Macau Special Administrative Region of China.

A warranty period of 3 years on parts and 1 year on labor means that IBM provides warranty service without charge for:

1. parts and labor during the first year of the warranty period; and
2. parts only, on an exchange basis, in the second and third years of the warranty period. IBM will charge you for any labor provided in performance of the repair or replacement(s) in the second and third year of the warranty period.

Machine type Topspin IB Host Channel Adapter Expansion Card for IBM eServer BladeCenter

Country or Region of Purchase	Warranty Period	Type of Warranty Service*
Worldwide	1 year	1
* See “Types of Warranty Service” for the legend and explanations of warranty-service types.		

Types of Warranty Service

If required, IBM provides repair or exchange service depending on the type of warranty service specified for your Machine in the above table and as described below. Warranty service may be provided by your reseller if approved by IBM to perform warranty service. Scheduling of service will depend upon the time of your call and is subject to parts availability. Service levels are response time objectives and are not guaranteed. The specified level of warranty service may not be available in all worldwide locations, additional charges may apply outside IBM’s normal service area, contact your local IBM representative or your reseller for country and location specific information.

Customer Replaceable Unit (“CRU”) Service

IBM provides replacement CRUs to you for you to install. CRU information and replacement instructions are shipped with your Machine and are available from IBM at any time on your request. Installation of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. You may install a Tier 2 CRU yourself or request IBM to install it, at no additional charge, under the type of warranty service designated for your Machine. IBM specifies in the materials shipped with a replacement CRU whether a defective CRU must be returned to IBM. When return is required, 1) return instructions and a container are shipped with the replacement CRU, and 2) you may be charged for the replacement CRU if IBM does not receive the defective CRU within 30 days of your receipt of the replacement.

3. On-site Service

IBM or your reseller will either repair or exchange the failing Machine at your location and verify its operation. You must provide suitable working area to allow disassembly and reassembly of the IBM Machine. The area must be clean, well lit and suitable for the purpose. For some Machines, certain repairs may require sending the Machine to an IBM service center.

4. Courier or Depot Service*

You will disconnect the failing Machine for collection arranged by IBM. IBM will provide you with a shipping container for you to return your Machine to a designated service center. A courier will pick up your Machine and deliver it to the designated service center. Following its repair or exchange, IBM will arrange the return delivery of the Machine to your location. You are responsible for its installation and verification.

5. Customer Carry-In or Mail-In Service

You will deliver or mail as IBM specifies (prepaid unless IBM specifies otherwise) the failing Machine suitably packaged to a location IBM designates. After IBM has repaired or exchanged the Machine, IBM will make it available for your collection or, for Mail-in Service, IBM will return it

to you at IBM's expense, unless IBM specifies otherwise. You are responsible for the subsequent installation and verification of the Machine.

6. **CRU and On-site Service**

This type of Warranty Service is a combination of Type 1 and Type 2 (see above).

7. **CRU and Courier or Depot Service**

This type of Warranty Service is a combination of Type 1 and Type 3 (see above).

8. **CRU and Customer Carry-In or Mail-In Service**

This type of Warranty Service is a combination of Type 1 and Type 4 (see above).

When a 5, 6 or 7 type of warranty service is listed, IBM will determine which type of warranty service is appropriate for the repair.

* This type of service is called ThinkPad® EasyServ or EasyServ in some countries.

The IBM Machine Warranty World Wide Web site at

http://www.ibm.com/servers/support/machine_warranties/ provides a worldwide overview of IBM's Limited Warranty for Machines, a Glossary of IBM definitions, Frequently Asked Questions (FAQs) and Support by Product (Machine) with links to Product Support pages. **The IBM Statement of Limited Warranty is also available on this site in 29 languages.**

To obtain warranty service contact IBM or your IBM reseller. In Canada or the United States, call 1-800-IBM-SERV (426-7378). In the EU countries, see the telephone numbers below.

EU Country Telephone List

Phone numbers are subject to change without notice. For the warranty service contact telephone number in a country subsequently added to the EU and not yet reflected in the list below, contact IBM in that country or visit the Web site above for a current telephone listing.

Austria -- +43-1-24592-5901	Latvia -- +386-61-1796-699
Belgium -- +32-70-23-3392	Lithuania -- +386-61-1796-699
Cyprus -- +357-22-841100	Luxembourg -- +352-298-977-5063
Czech Republic -- +420-2-7213-1316	Malta -- +356-23-4175
Denmark -- +45-4520-8200	Netherlands -- +31-20-514-5770
Estonia -- +386-61-1796-699	Poland -- +48-22-878-6999
Finland -- +358-8001-4260	Portugal -- +351-21-892-7147
France -- +33-238-557-450	Slovakia -- +421-2-4954-1217
Germany -- +49-1805-253553	Slovenia -- +386-1-4796-699
Greece -- +30-210-680-1700	Spain -- +34-91-714-7983
Hungary -- +36-1-382-5720	Sweden -- +46-8-477-4420
Ireland -- +353-1-815-4000	United Kingdom -- +44-1475-555-055
Italy -- +39-800-820-094	

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing

IBM Corporation

North Castle Drive

Armonk, NY 10504-1785

U.S.A.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Edition notice

© Copyright International Business Machines Corporation 2004. All rights reserved.

U.S. Government Users Restricted Rights — Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

Active Memory	Predictive Failure Analysis
Active PCI	PS/2
Active PCI-X	ServeRAID
Alert on LAN	ServerGuide
BladeCenter	ServerProven
C2T Interconnect	TechConnect
Chipkill	ThinkPad
EtherJet	Tivoli
e-business logo	Tivoli Enterprise
Eserver	Update Connector
FlashCopy	Wake on LAN
IBM	XA-32
IBM (logo)	XA-64
IntelliStation	X-Architecture
NetBAY	XcelL4
Netfinity	XpandOnDemand
NetView	xSeries
OS/2 WARP	

Intel, MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Adaptec and HostRAID are trademarks of Adaptec, Inc., in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Red Hat, the Red Hat “Shadow Man” logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Important notes

Processor speeds indicate the internal clock speed of the microprocessor; other factors also affect application performance.

CD-ROM drive speeds list the variable read rate. Actual speeds vary and are often less than the maximum possible.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1 000 000 bytes, and GB stands for approximately 1 000 000 000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Maximum memory may require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven®, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Product recycling and disposal

This unit contains materials such as circuit boards, cables, electromagnetic compatibility gaskets, and connectors which may contain lead and copper/beryllium alloys that require special handling and disposal at end of life. Before this unit is disposed of, these materials must be removed and recycled or discarded according to applicable regulations. IBM offers product-return programs in several countries. Information on product recycling offerings can be found on IBM’s Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.

Battery return program

This product may contain a sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/batteryrecycle.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a collection process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and battery packs from IBM

equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Have the IBM part number listed on the battery available prior to your call.

In the Netherlands, the following applies.



廢電池請回收

Electronic emission notices

Federal Communications Commission (FCC) statement



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 not installed and used in accordance with the instruction manual, 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.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

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, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Class A emission compliance statement

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

Avis de conformité à la réglementation d'Industrie Canada

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

Australia and New Zealand Class A statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

United Kingdom telecommunications safety requirement

Notice to Customers

This apparatus is approved under approval number NS/G/1234/J/100003 for indirect connection to public telecommunication systems in the United Kingdom.

European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Taiwanese Class A warning statement

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。

Chinese Class A warning statement

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。

Japanese Voluntary Control Council for Interference (VCCI) statement

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。

Safety

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية.

Antes de instalar este produto, leia as Informações de Segurança.

在安装本产品之前，请仔细阅读 **Safety Information** (安全信息)。

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Ennen kuin asennat tämän tuotteen, lue turvaohjeet kohdasta Safety Information.

Avant d'installer ce produit, lisez les consignes de sécurité.

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφαλείας (safety information).

לפני שתתקינו מתצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

製品の設置の前に、安全情報をお読みください。

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Przed zainstalowaniem tego produktu, należy zapoznać się z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Antes de instalar este produto, leia as Informações sobre Segurança.

Перед установкой продукта прочтите инструкции по технике безопасности.

Pred inštaláciou tohto zariadenia si pečítajte Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Antes de instalar este producto, lea la información de seguridad.

Läs säkerhetsinformationen innan du installerar den här produkten.

Statement 1:**DANGER**

Electrical current from power, telephone, and communication cables is hazardous.

To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

Statement 3:

**CAUTION:**

When laser products (such as CD-ROMs, DVD drives, fiber optic devices, or transmitters) are installed, note the following:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

**DANGER**

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following.

Laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam.

Class 1 Laser Product
Laser Klasse 1
Laser Class 1
Luokan 1 Laserlaitte
Appareil À Laser de Classe 1

Statement 21:

**CAUTION:**

Hazardous energy is present when the blade is connected to the power source. Always replace the blade cover before installing the blade.

About the Topspin BladeCenter HCA

The following sections appear in this chapter:

- [“Introduction” on page 1](#)
- [“HCA Expansion Cards and Server Switch Modules” on page 2](#)
- [“Supported Protocols” on page 3](#)

Introduction

The Topspin IB Host Channel Adapter Expansion Card for IBM eServer BladeCenter (hereafter “HCA Expansion Card”) provides InfiniBand I/O capability to processor blades in the IBM BladeCenter unit. The HCA adds two InfiniBand ports to the CPU blade cards to create an IB-capable cluster-in-a-box. [Figure 1-1](#) displays the HCA Expansion Card.



Figure 1-1: Infiniband Host Channel Adapter Expansion Card for IBM eServer BladeCenter

The HCA Expansion Cards communicate with one another through the InfiniBand Server Switch Module for IBM eServer BladeCenter. For information on the Server Switch Module, refer to the *InfiniBand Switch Module for IBM eServer BladeCenter User Guide* that arrives with the Server Switch Module.

Features

- 133 MHz PCI-X Interface to dual 1x Infiniband bridge
- Two 1x Infiniband interfaces, line rate of the interfaces are 2.5Gbps per link theoretical maximum
- 128M Byte Table Memory (133MHz DDR SDRAM)
- I2C Serial EEPROM holding system Vital Product Data (VPD)
- IBM's proprietary Blade Daughter card form factor
- Identical operation and interfaces for existing Topspin's HCA drivers
- Forced air cooling compatible for highly reliable operation
- Port-to-port failover for select protocols (see [“Supported Protocols”](#) for details)

Supported Blade Servers

lists the IBM blade servers that support the HCA Expansion Card, and includes the maximum number of HCA Expansion Cards that the blade supports.

Table 1-1: Supported Blade Servers

Blade Server	Maximum Number of Supported HCA Expansion Cards
HS20 8678	1
HS20 8832	1
HS40 8839	2
HS20 8843	2
JS20	1

HCA Expansion Cards and Server Switch Modules

Within the BladeCenter chassis, Server Switch Modules manage traffic to and from HCA Expansion Cards on the BladeCenter hosts. Each HCA expansion card adds two IB ports to a BladeCenter host. Each HCA port connects via the BladeCenter unit backplane to a particular Server Switch Module bay. The first IB port of each HCA card (ib0) connects to the Server Switch Module in bay 3, and the second IB port of each HCA card (ib1) connects to bay 4. With the Server Switch Module and HCA Expansion Cards, you can create a non-redundant, single-switch topology or a redundant, dual-switch topology.

Single-Switch Topology

When you populate just one BladeCenter module bay with an Server Switch Module, you create a bi-sectional bandwidth topology. However, this topology does *not* provide redundant links from the HCA Expansion Cards to the Server Switch Module. We strongly recommend that you implement a [“Dual-Switch Topology”](#) to avoid single points of failure.

Dual-Switch Topology

To enable IB redundancy on the BladeCenter unit, you must install one Server Switch Module in each available bay. HCA Expansion Cards do not support redundant links to a single switch bay. When you add a second Server Switch Module to the BladeCenter unit, each port of each HCA Expansion Card connects to one Server Switch Module.

Supported Protocols

The Topspin HCA expansion module supports the following protocols:

- Internet Protocol over InfiniBand (IPoIB)
- Socket Direct Protocol (SDP)
- User Direct Access Programming Library (uDAPL)
- SCSI RDMA Protocol (SRP)
- Message Passing Interface (MPI)



NOTE: Topspin supports all protocols for Linux. As of this release, Topspin only supports IPoIB and SRP for Windows.

IPoIB

IPoIB allows IP networks to utilize the InfiniBand fabric. SDP and uDAPL use IPoIB to resolve IP addresses. You configure IPoIB as you would a standard Ethernet interface. IPoIB automatically adds IB interface names to the network configuration. The interface names correspond to the ports on the HCA. For configuration instructions, refer to [“Configure IPoIB” on page 35](#).



NOTE: IPoIB supports port-to-port failover on the HCA daughter card.

SDP

SDP provides a high-performance, zero-copy data transfer protocol used for stream-socket networking over an InfiniBand fabric. You can configure the Topspin driver to automatically translate TCP to SDP based on source IP, destination, or program name. For configuration instructions, refer to [“Configure SRP” on page 45](#).

uDAPL

The User Direct Access Programming Library (uDAPL) defines a set of APIs that leverages InfiniBand’s remote direct memory access (RDMA) capabilities. uDAPL is installed transparently with the Topspin driver library. Your application must explicitly support uDAPL. This library requires no manual configuration. However, if your application supports uDAPL, it may require additional configuration changes. Please refer to your application documentation for more information.

SRP

SRP runs SCSI commands across RDMA-capable networks so InfiniBand hosts can communicate with Fibre Channel storage devices. This information is used to assign devices and mount file systems so that the data on those file systems is accessible to the host. The SRP driver is installed as part of the Topspin driver package and is loaded automatically upon host reboot. This protocol requires that a Fibre Channel gateway be present in the IB fabric. For configuration instructions, refer to [“Configure SRP” on page 45](#).



NOTE: SRP supports port-to-port failover on the HCA daughter card.

MPI

MPI source is included in the installation package. MPI is also available on the Topspin support Web site (<http://support.topspin.com>, **Software Downloads** link).



NOTE: MPI software does not appear on your installation CD. You must download the software from the support Web site at <http://support.topspin.com>.

- There is no restriction on which HCA port is used.
- Support for Opteron 64 bit operation is provided.

For configuration instructions, refer to “[Configure MPI \(Linux Only\)](#)” on page 47.

Installing the HCA Expansion Card

The following sections appear in this chapter:

- [“Install the HCA Expansion Card Hardware” on page 5](#)
- [“Prepare to Install Drivers” on page 17](#)
- [“Install Linux Drivers” on page 17](#)
- [“Install Windows 2000 Drivers” on page 17](#)
- [“Install Windows 2003 Drivers” on page 21](#)
- [“Upgrade Firmware” on page 25](#)
- [“Uninstall Windows Host Drivers \(2000 or 2003\)” on page 27](#)
- [“Configure High-Availability Ports on Linux Hosts” on page 27](#)
- [“Configure High-Availability Ports on Windows Hosts” on page 28](#)

Install the HCA Expansion Card Hardware

To install the Topspin BladeCenter HCA hardware into a BladeCenter unit, perform the following steps:

1. Ground yourself appropriately.
2. Identify the blade server in the chassis into which you want to insert the HCA.
3. Refer to the appropriate blade server documentation for instructions on how to install an expansion card.



NOTE: The HCA comes with a Bladeserver Unit Expansion Card tray. Refer to your IBM Bladeserver Unit documentation for installation instructions.

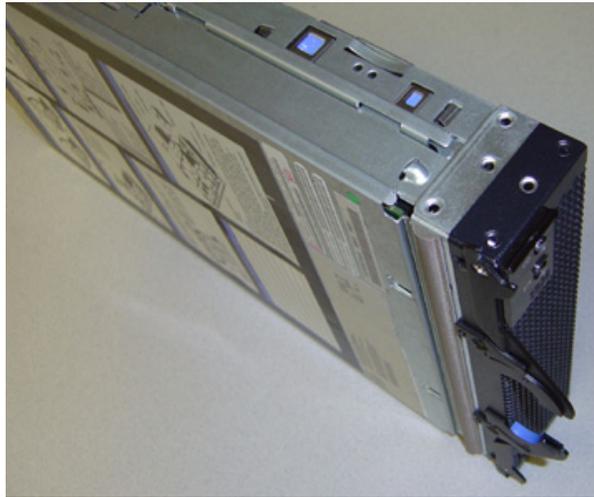
4. Proceed to [“Prepare to Install Drivers”](#).

Install Multiple HCAs in the HS40 8839 Blade

These instructions are supplemental to the blade server documentation. To install two HCA Expansion Cards in an HS40 blade, perform the following steps:

1. Remove your HS40 blade from your BladeCenter unit.

Example



2. Open the lid of the blade server.

Example



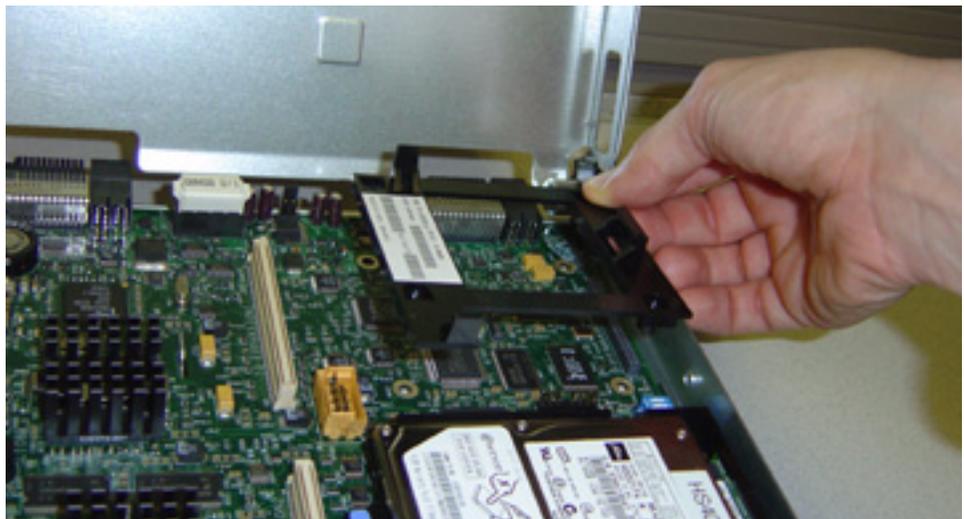
3. Remove the screws from the far HCA Expansion Card bay.

Example



4. Place the HCA tray in the blade server.

Example



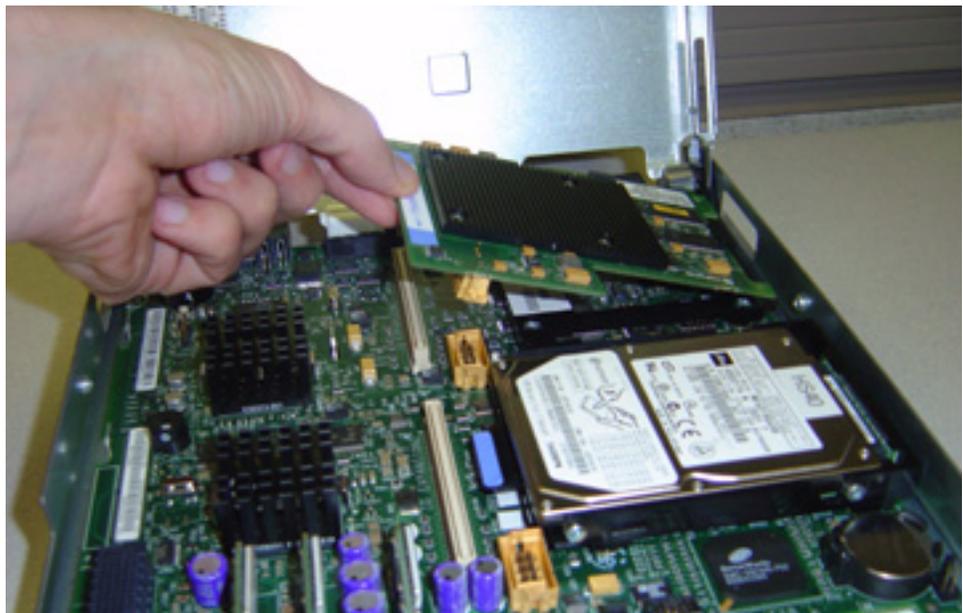
- Secure the tray to the blade server with the screws provided.

Example



- Insert the HCA Expansion Card into the tray.

Example



7. Press the HCA Expansion Card firmly into place.

Example



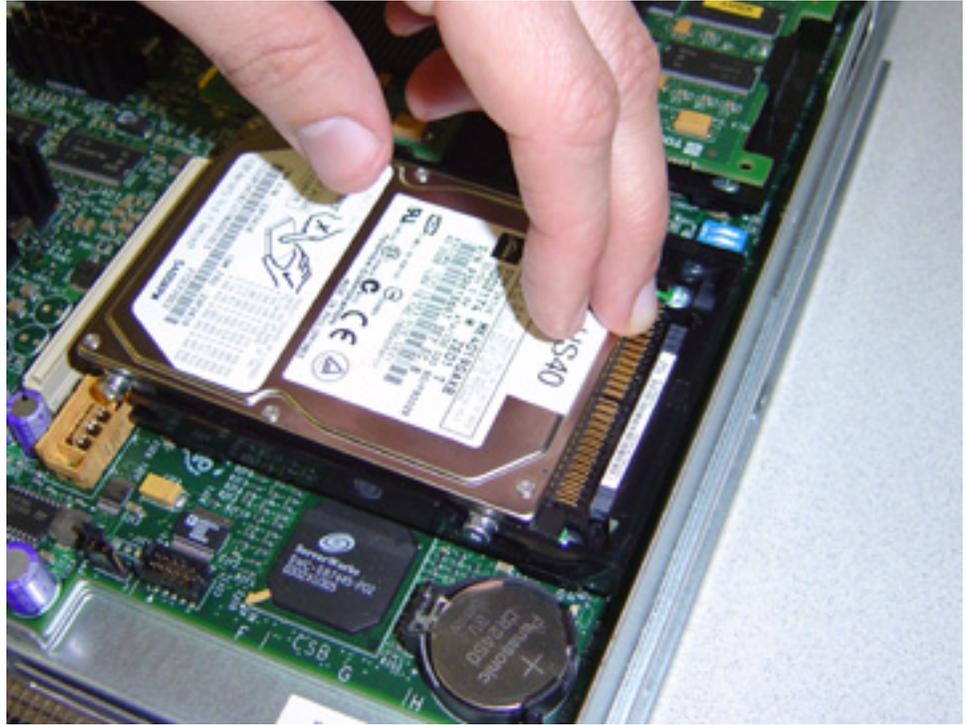
8. (If applicable) Remove the IDE hard drive.
 - a. Depress the tab on the hard disk tray.

Example



- b. Pull the hard disk out of the tray until the pins clear the tray.

Example



- c. Remove the hard disk from the blade server.

Example



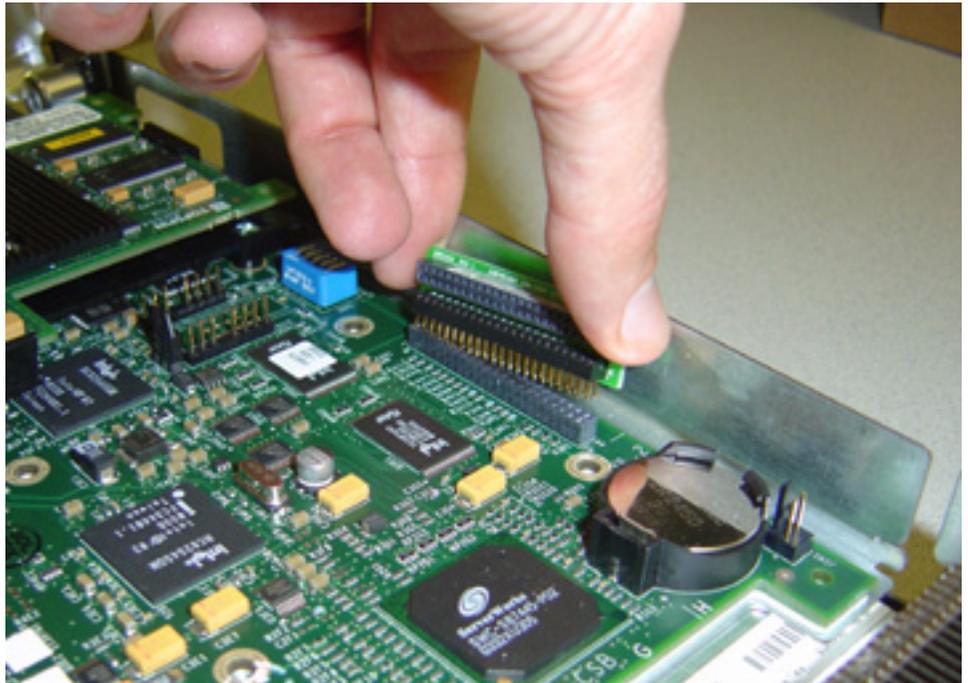
- d. Unscrew the hard disk tray and remove it from the blade server.

Example



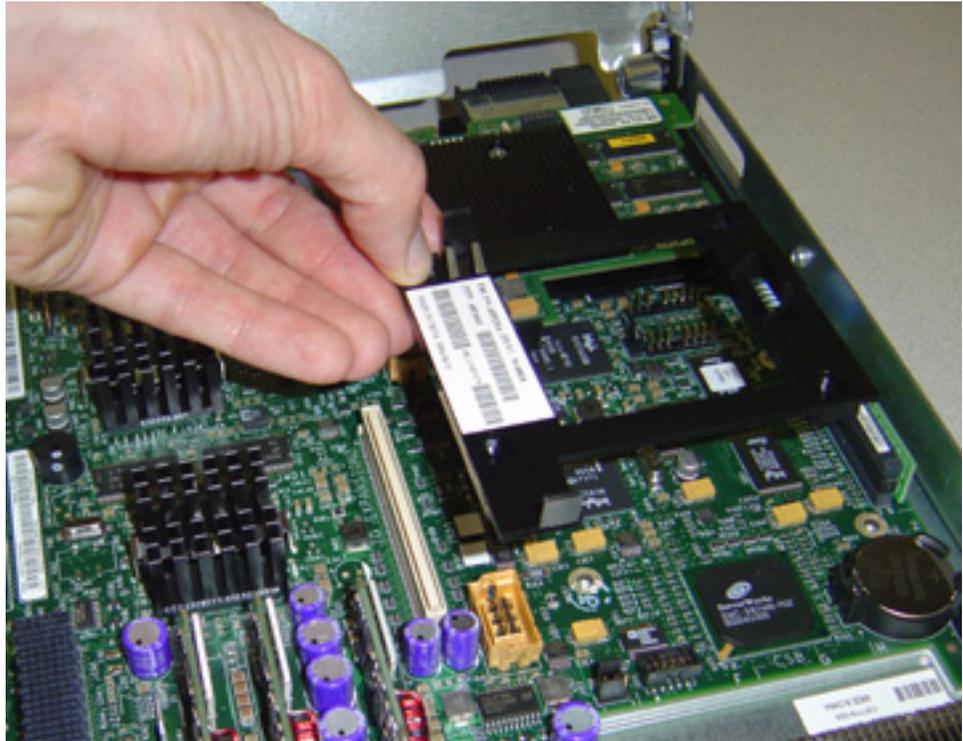
- e. Remove the IDE hard drive connector.

Example



9. Place the HCA Expansion Card tray in the blade server.

Example



10. Screw the HCA Expansion Card tray in the blade server and secure it with the screws provided.

Example



11. Insert the HCA Expansion Card into the tray.

Example

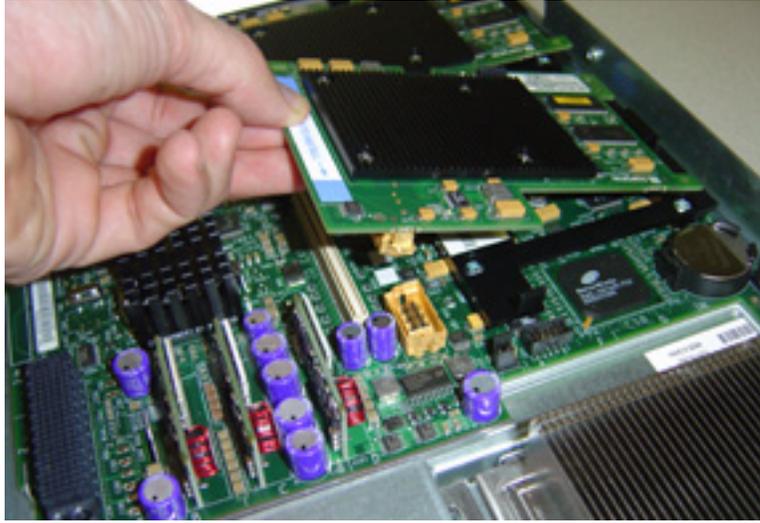


Figure 2-1 displays the HS40 blade with two HCA Expansion Cards installed.



Figure 2-1: HS40 Blade with 2 HCA Expansion Cards

Install Two HCAs in an HS20 8843 Blade with a Blade Server Extension

These instructions are supplemental to the blade server documentation. To install two HCA Expansion Cards in an HS20 blade with a blade server extension (BSE), perform the following steps:

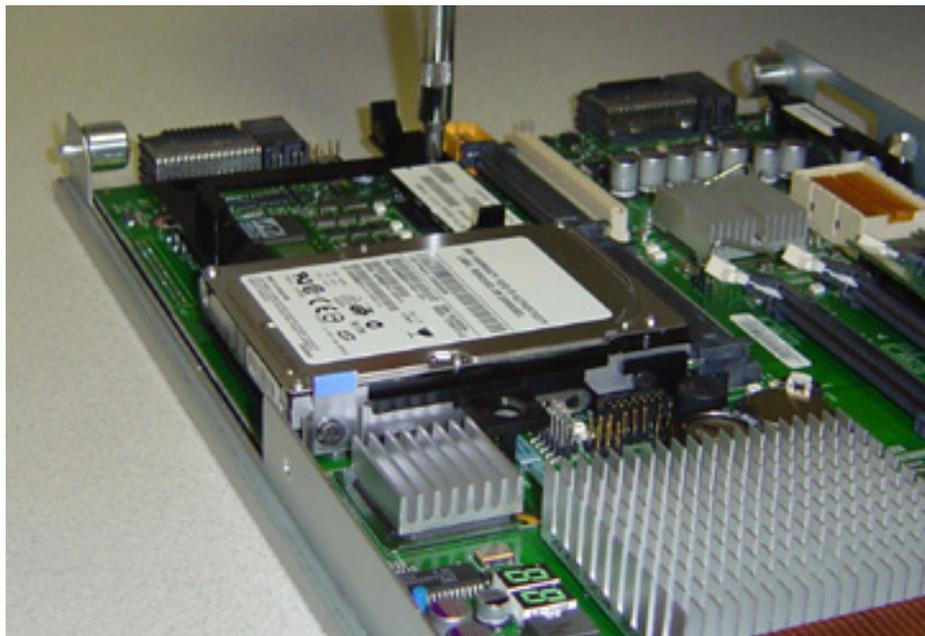
1. Place the HCA Expansion Card tray in the blade server.

Example



2. Secure the HCA Expansion Card tray to the blade server with the screws provided.

Example



5. Firmly seat the HCA Expansion Card in the tray.

Example

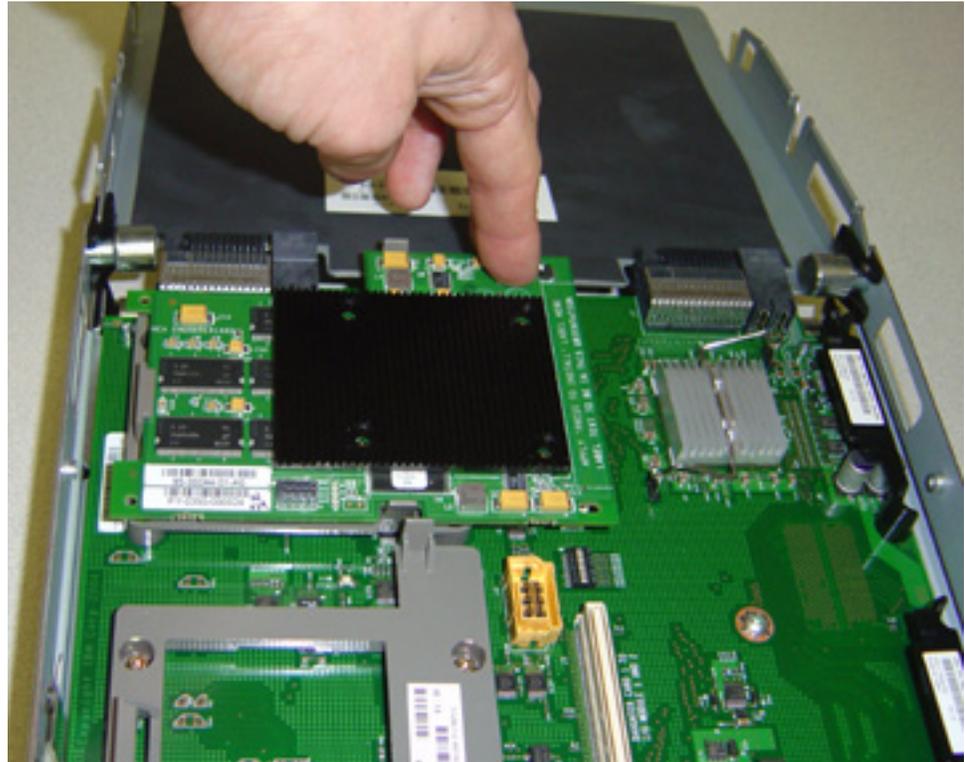


Figure 2-2 displays the HCA Expansion Card in the BSE.



Figure 2-2: HCA Expansion Card in the BSE

Prepare to Install Drivers

The HCA Expansion Card supports Linux and Windows drivers. To install drivers on a BladeCenter host, press the CD and monitor buttons on the appropriate Blade Server host (for more information, refer to your BladeCenter unit documentation). The host gains access to the CD-ROM drive, mouse, keyboard, and monitor.

Install Linux Drivers

To install Linux drivers, perform the following steps:

1. Log in to your Linux host.
2. Place the Topspin host driver CD-ROM in the CD-ROM drive.
3. Mount your CD-ROM drive (if required).

Example

```
# mount /media/cdrom/
```

4. Move to the top-level directory of the CD-ROM.

Example

```
# cd /media/cdrom/
```

5. (Optional) Enter the **rpm -qa | grep topspin** command to view the Topspin installation files.

```
qa-bc1a-blade2:~ # rpm -qa | grep topspin
topspin-ib-mod-sles9-2.6.5-7.97smp-3.0.0-179
topspin-ib-sles9-3.0.0-179
qa-bc1a-blade2:~ #
```

6. Enter the **tsinstall** command to install the host drivers. This script automatically detects the available kernel and installs the appropriate RPM packages.

Example

```
host:/media/cdrom # ./tsinstall
```

```
The following kernels are installed, but do not have drivers available:
 2.6.5-7.97smp.x86_64
```

```
The following drivers are currently installed, but do not have new replacement
drivers:
```

```
  topspin-ib-mod-sles9-2.6.5-7.97smp-3.0.0-179.x86_64
```

```
Would you like to remove these drivers? (y/n) n
```

```
host:/media/cdrom #
```

Install Windows 2000 Drivers

To install Windows drivers, determine your installation type, determine your OS, and execute the appropriate installation steps. You may perform any of the following installation types:

- IPoIB and SRP drivers (.EXE file).

- IPoIB and SRP drivers (.EXE file) with the Boot over InfiniBand feature.
- WHQL drivers installation (.ZIP file).



NOTE: This installation meets Windows Hardware Quality Labs (WHQL) requirements. Initial WHQL installations omit certain utilities, but you can add those utilities with an optional second step. For details, see [“Install WHQL Drivers on a Windows 2000 or 2003 Server \(both BoIB and non-BoIB\)”](#) on page 25.

You may install host drivers on either of the following Windows operating systems:

- Windows 2000
- Windows 2003

Install Non-WHQL Drivers on a Windows 2000 Server (both BoIB and non-BoIB)

To install Windows drivers after you have installed the HCA Expansion Card, perform the following steps:

1. Log in to your Windows 2000 host. A **Found New Hardware** alert appears ([Figure 2-3](#)).

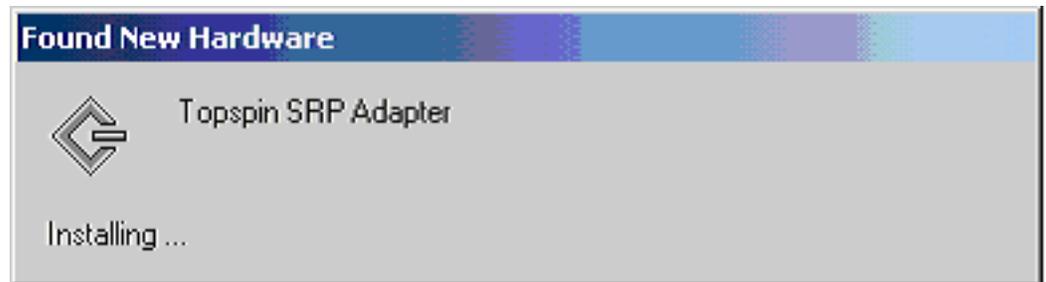


Figure 2-3: Found New Hardware Alert.

2. Ignore the alert. The **Found New Hardware Wizard** opens (Figure 2-4). Click **Cancel**.



Figure 2-4: Found New Hardware Wizard



NOTE: This wizard may not appear immediately. When the wizard does appear, click **Cancel**.

3. Insert your driver CD (or download drivers from <http://support.topspin.com>) and run `Topspin-ib-W2k-releaseNumber.exe` from your CD. The **Topspin InfiniBand SDK Setup** window opens (Figure 2-5).

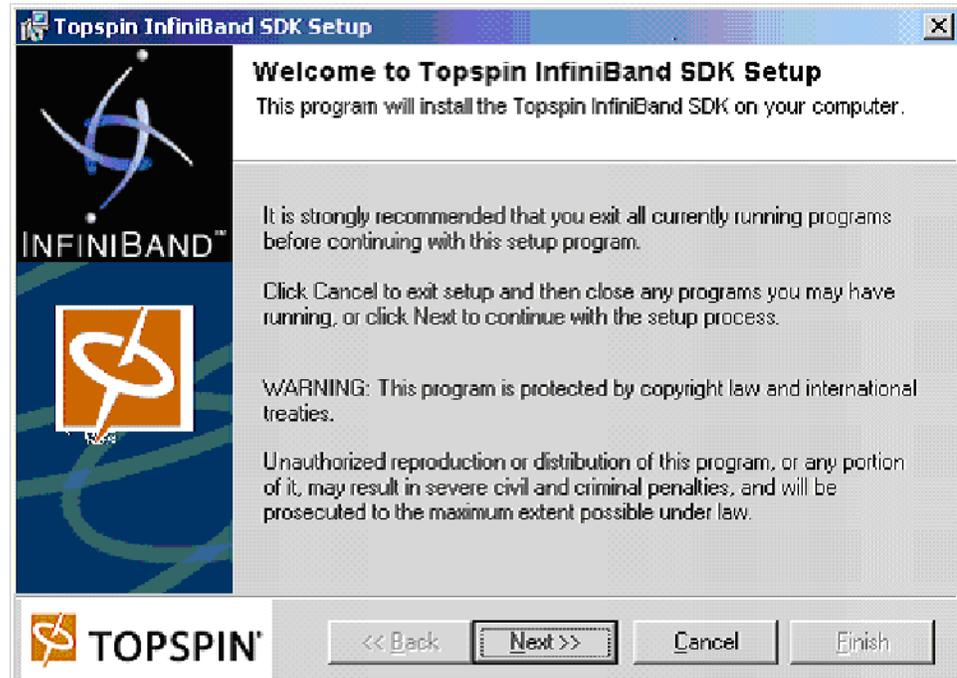


Figure 2-5: Welcome Screen

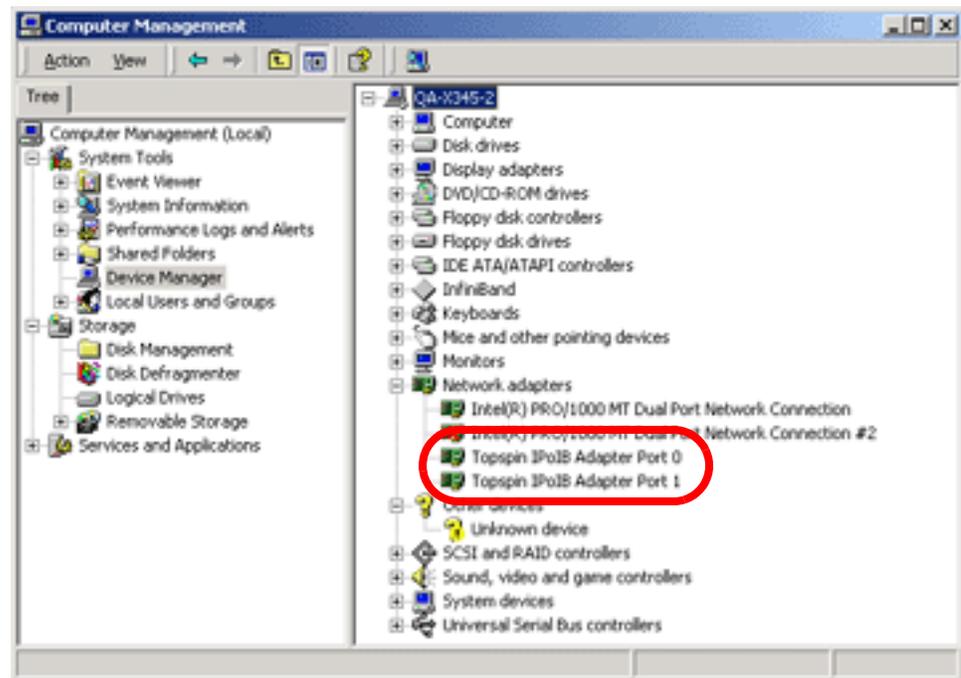
4. Click **Next**. The **End User License Agreement** opens.
5. Click the **I Agree to the terms listed above** radio button, then click **Next**. The **Product Information** screen appears.
6. Click **Next**. The **Ready to Install** screen opens.
7. Click **Next**. The installation executes. The Installation Complete screen appears when the installation completes.
8. Click **Finish**. The system prompts you to reboot. If you have not previously installed host drivers on this system, you do not need to reboot.
9. Click **Yes**. Your host reboots.
10. Proceed to [“Verify the Installation” on page 20](#) or [“Configure IPoIB on Windows” on page 41](#).

Verify the Installation

1. Log in to your host.
2. Click **Start > Programs > Topspin InfiniBand SDK**. If icons appear in the **Topspin InfiniBand SDK** folder, the installation succeeded.
3. Click **Start > Settings > Network and Dial-up Connections**. The **Network and Dial-up Connections** window opens.
4. Verify that a **Topspin IP over IB Virtual Channel Adapter** entry appears for each HCA *port*.
5. Right-click the **My Computer** icon on your desktop, then select **Manage** from the right-click menu. The **Computer Management** window opens.

- Click **Device Manager**, then expand **Network Adapters** to verify your IB interfaces function properly. Verify that both ports appear in the display, with no error icons (see example).

Example



Install Windows 2003 Drivers

To install drivers after you install the HCA, perform the following steps:

- Log in to your Windows 2003 host. Insert the driver CD or download the Win2k3 drivers.
- Begin the installation. The **Product Install** window appears.



Figure 2-6: Begin Product Installation

3. Click the **Next** button. The End User License Agreement appears.
4. Click the **I Agree to the terms listed above** radio button and click the **Next** button to continue. The Product Information window appears.
5. Click the **Next** button. The **Installation in Progress** window appears with the status of the installation.

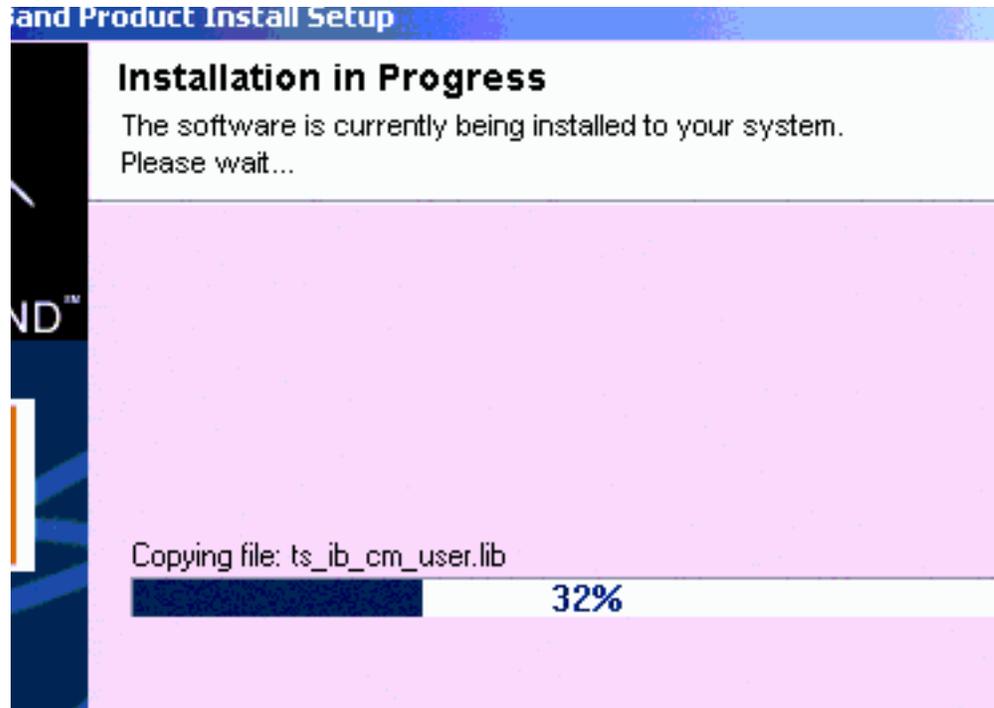


Figure 2-7: View the Installation Status

A Windows driver dialog box appears. This dialog box explains that the new drivers do not contain a Windows signature. Although the drivers are not digitally signed, they are tested using the Microsoft HCT suite. WHQL submission materials are available per request (CAB files, etc.).

6. Click the **Continue Anyway** button each time the Windows Logo testing dialog box appears.

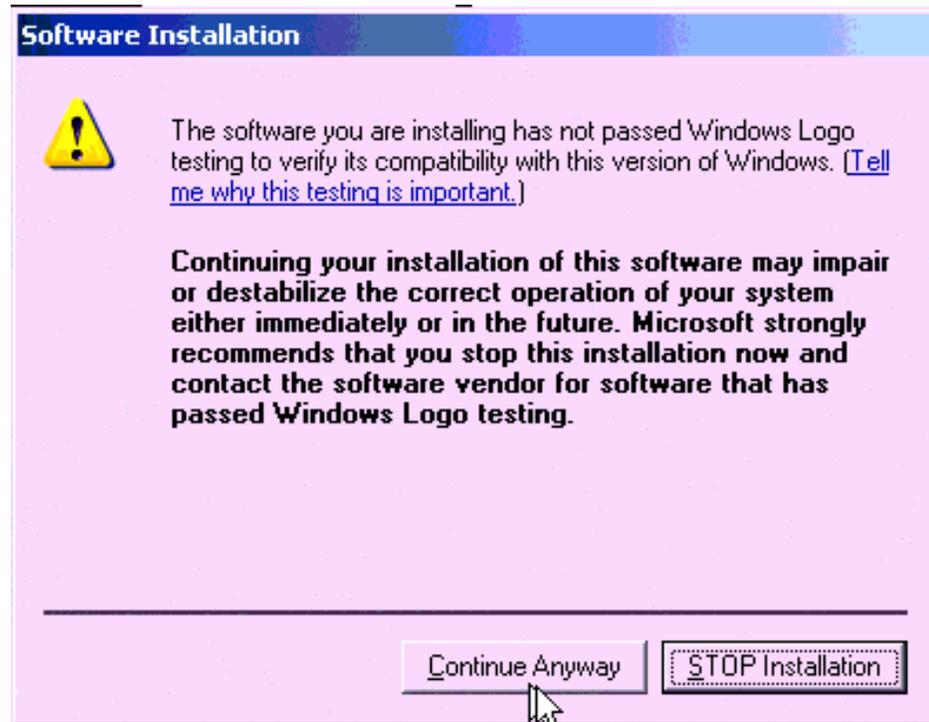


Figure 2-8: Click Continue Anyway

The following Windows dialog box appears for each HCA you have installed.

7. Click the **Yes** button.



Figure 2-9: Click the Yes Button

8. **Important:** Ignore the **Found New Hardware Wizard** when it appears, as it will continue without user intervention. This window appears three times (twice for the IPoIB and once for SRP).

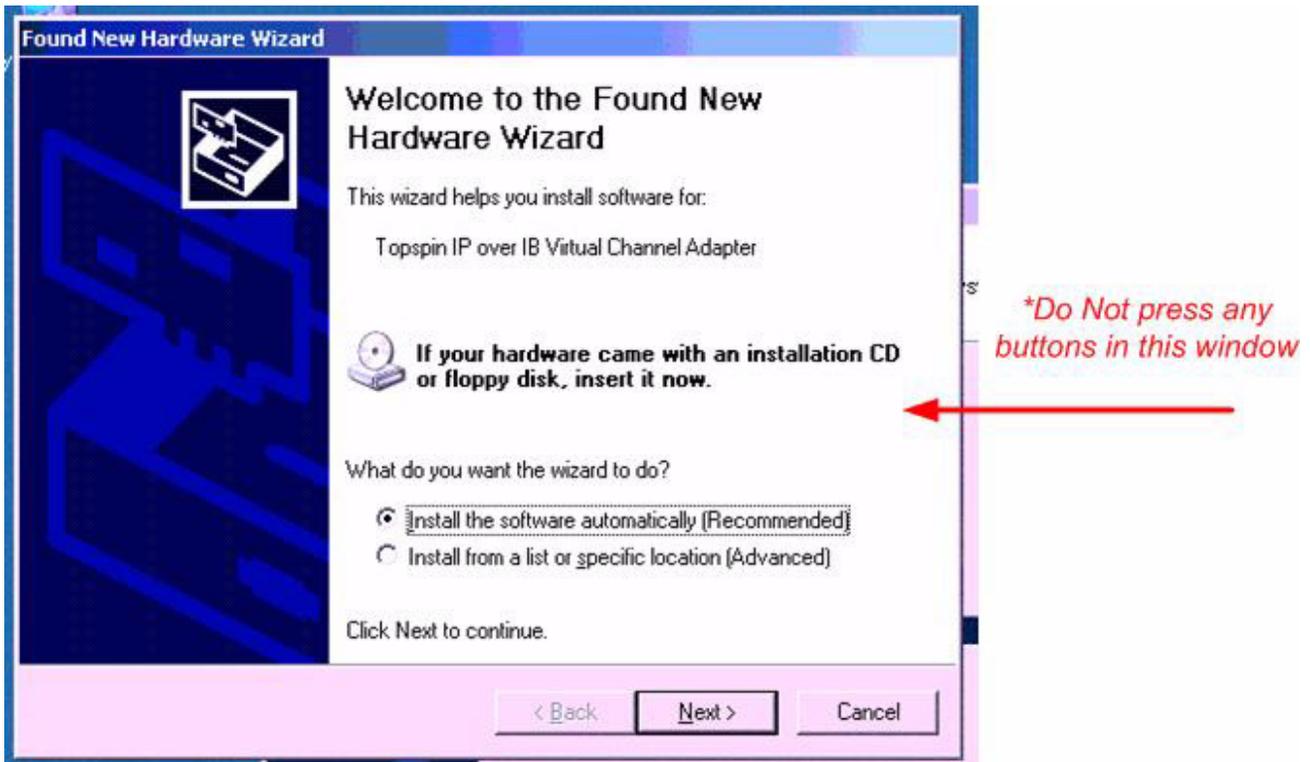


Figure 2-10: Ignore the Found New Hardware Wizard Windows

The **Installation Complete** window appears.

9. Click the **Finish** button.

10. Reboot the server.
 - a. Select **Start > Shutdown**.
 - b. Select **Restart** from the drop-down menu.
 - c. (Optional) Enter a reason for the system reboot in the **Comment** field, if your server requires it.



Figure 2-11: Click the Yes Button

- d. Click **OK**. The system reboots and new driver installation is complete.

Install WHQL Drivers on a Windows 2000 or 2003 Server (both BoIB and non-BoIB)

To perform a WHQL installation, refer to the documentation at <http://support.topspin.com>.

Upgrade Firmware

When you receive a new Windows driver or Linux RPM, perform the procedures in this section to upgrade your firmware.



NOTE: When you upgrade host drivers, the process upgrades firmware automatically.

Upgrade Firmware with Linux

To update Linux drivers, perform the following steps:

1. Log in to your Linux host.
2. Place the Topspin host driver CD-ROM in the CD-ROM drive.

3. Mount your CD-ROM drive (if required).

Example

```
# mount /mnt/cdrom/
```

4. Move to the top-level directory of the CD-ROM.

Example

```
# cd /mnt/cdrom/
```

5. Enter the **tsinstall** command to install the host drivers. This script automatically detects the available kernel and overwrites the existing firmware with the appropriate RPM packages.

Example

```
# ./tsinstall
```

```
The following kernels are installed, but do not have drivers available:
 2.4.20-8.i686
```

```
The following installed packages are out of date and will be upgraded:
topspin-ib-mod-rh9-2.4.20-8smp-1.1.3-666.i686
```

```
The following packages will be installed:
topspin-ib-rh9-1.1.3-687.i686.rpm (libraries, binaries, etc)
topspin-ib-mod-rh9-2.4.20-8smp-1.1.3-687.i686 (drivers)
```

```
installing 100%
```

```
#####
#####
```

Upgrade Firmware with Windows

To upgrade your Windows drivers, perform the following steps:

1. Log in to your Windows host and place the firmware CD in the BladeCenter unit CD-ROM drive.
2. Click **Start > Programs > Topspin InfiniBand SDK > Utilities**. The **Utilities** window opens and displays a list of utilities.
3. Enter the **tvflash** command and press **Enter** to view the tvflash options and descriptions.
4. (Optional) enter **tvflash -s** and *filename.bin* to save your existing firmware version to the **util** directory.
5. Enter the **tvflash -i** command to display the type of the HCA and the firmware version that the HCA runs.
6. Browse your CD for your firmware files and select a firmware file whose type (Cougar, Cougar Cub, or Jaguar) matches your HCA type.
7. Enter the **tvflash -d [path/filename]** to install the firmware on the HCA.
8. Enter **tvflash -i** to verify the new firmware version.

Uninstall Windows Host Drivers (2000 or 2003)

To uninstall drivers on Windows machines, perform the following steps:

1. Log in to your Windows host.
2. Click **Start > Programs > Topspin InfiniBand SDK > Uninstall**. The **Product Uninstall** window opens.
3. Click the **Next** button. The **Removal Complete** window appears.
4. Click the **Finish** button.
5. Reboot your host.

Configure High-Availability Ports on Linux Hosts

IPoIB supports active/passive port failover high availability (HA) between two ports on the same HCA. When you enable the HA feature, both ports on the HCA (ib0 and ib1) merge into one virtual port. If you configure high availability between the ports on the HCA, only one of the physical ports passes traffic. The other port is used as standby in the event of a failure.

Merge Physical Ports

To configure high availability on HCA ports in a Linux host, perform the following steps:

1. Show the available interfaces. This example shows two HCAs in a single host.

Example

```
[root@server-1 root]# /usr/local/topspin/sbin/
[root@server-1 sbin]# ./ipoibcfg list
ib0 (P_Key 0xffff)
ib1 (P_Key 0xffff)
ib2 (P_Key 0xffff)
ib3 (P_Key 0xffff)
```

2. Disable the interfaces that you want to merge into one virtual port.

Example

```
[root@server-1 sbin]# ifconfig ib0 down
[root@server-1 sbin]# ifconfig ib1 down
```

3. Enter the **ipoibcfg merge <interface> <interface>** command.

In the following example, **ipoibcfg merge ib0 ib1** merges the two ports on the first HCA into one virtual HA port.

Example

```
[root@server-1 sbin]# ./ipoibcfg merge ib0 ib1
```

- Show the interfaces. Note that ib1 no longer appears, as it is merged with ib0.

Example

```
[root@server-1 sbin]# ./ipoibcfg list
ib0 (P_Key 0xffff)
ib2 (P_Key 0xffff)
ib3 (P_Key 0xffff)
[root@server-1 sbin]# ifconfig ib0 down
[root@server-1 sbin]# ifconfig ib1 down
```

- Re-enable the interface.

Example

```
[root@server-1 sbin]# ifconfig ib0 up
```

- Assign an IP address to the merged port just as you would to a standard interface. Refer to [“Configure IPoIB on Linux” on page 36](#).

Unmerge Physical Ports

To disable active/passive high-availability on the physical ports, perform the following steps:

- Disable the interface that you want to unmerge.

Example

```
[root@server-1 sbin]# ifconfig ib0 down
```

- Enter the `ipoibcfg unmerge <interface>` command.

Example

```
[root@server-1 sbin]# ./ipoibcfg unmerge ib0
```

- Re-enable the interfaces.

Example

```
[root@server-1 sbin]# ifconfig ib0 up
[root@server-1 sbin]# ifconfig ib1 up
```

Configure High-Availability Ports on Windows Hosts

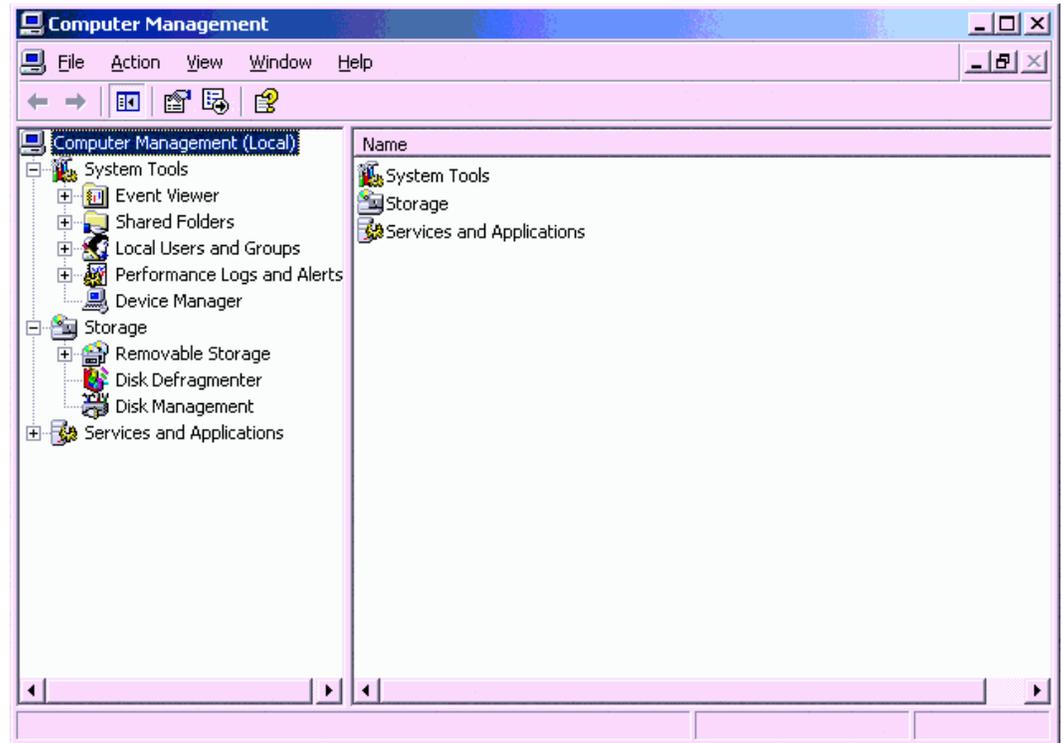
The two ports on a single HCA can both be used as active ports, or the two ports can be merged into an active/passive configuration to allow for failover between the ports.

By default, both ports are defined as separate, active ports and fail-over must be explicitly configured.

To create a single “logical” interface for two physical ports on a single HCA, perform the following steps:

1. Remove any previously assigned IP address for ports ib0 and ib1.
2. Navigate to your desktop, right-click the **My Computer** icon, and select **Manage** from the right-click menu. The **Computer Management** window opens.

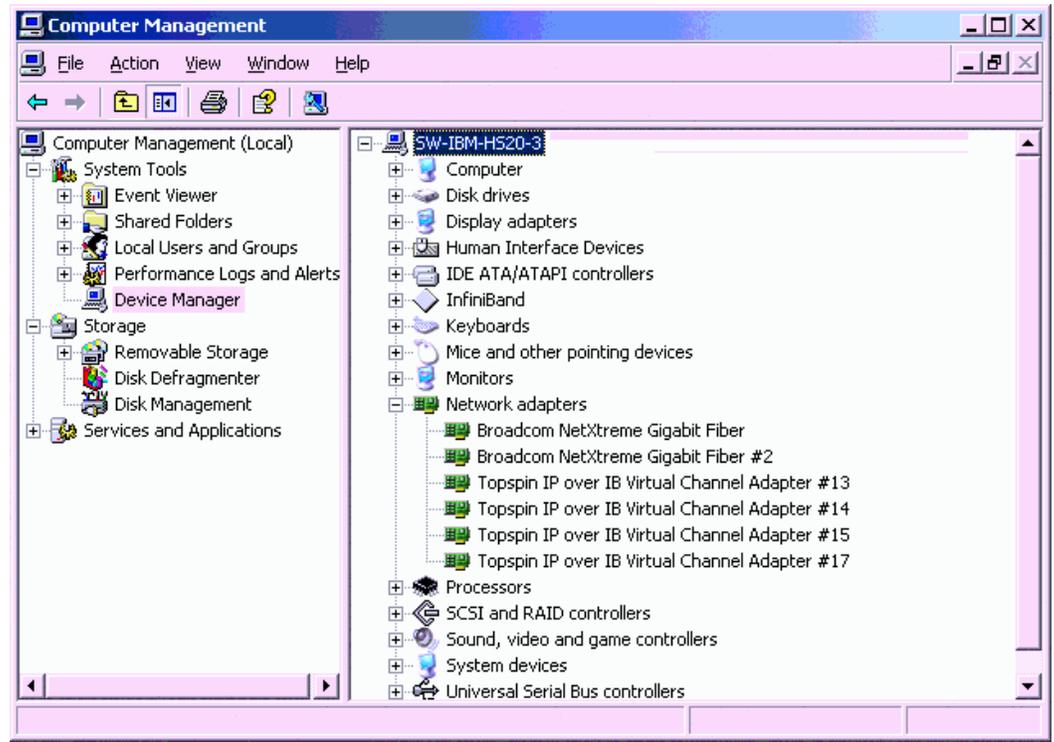
Example



3. Click **Device Manager** in the left-hand frame of the **Computer Management** window.

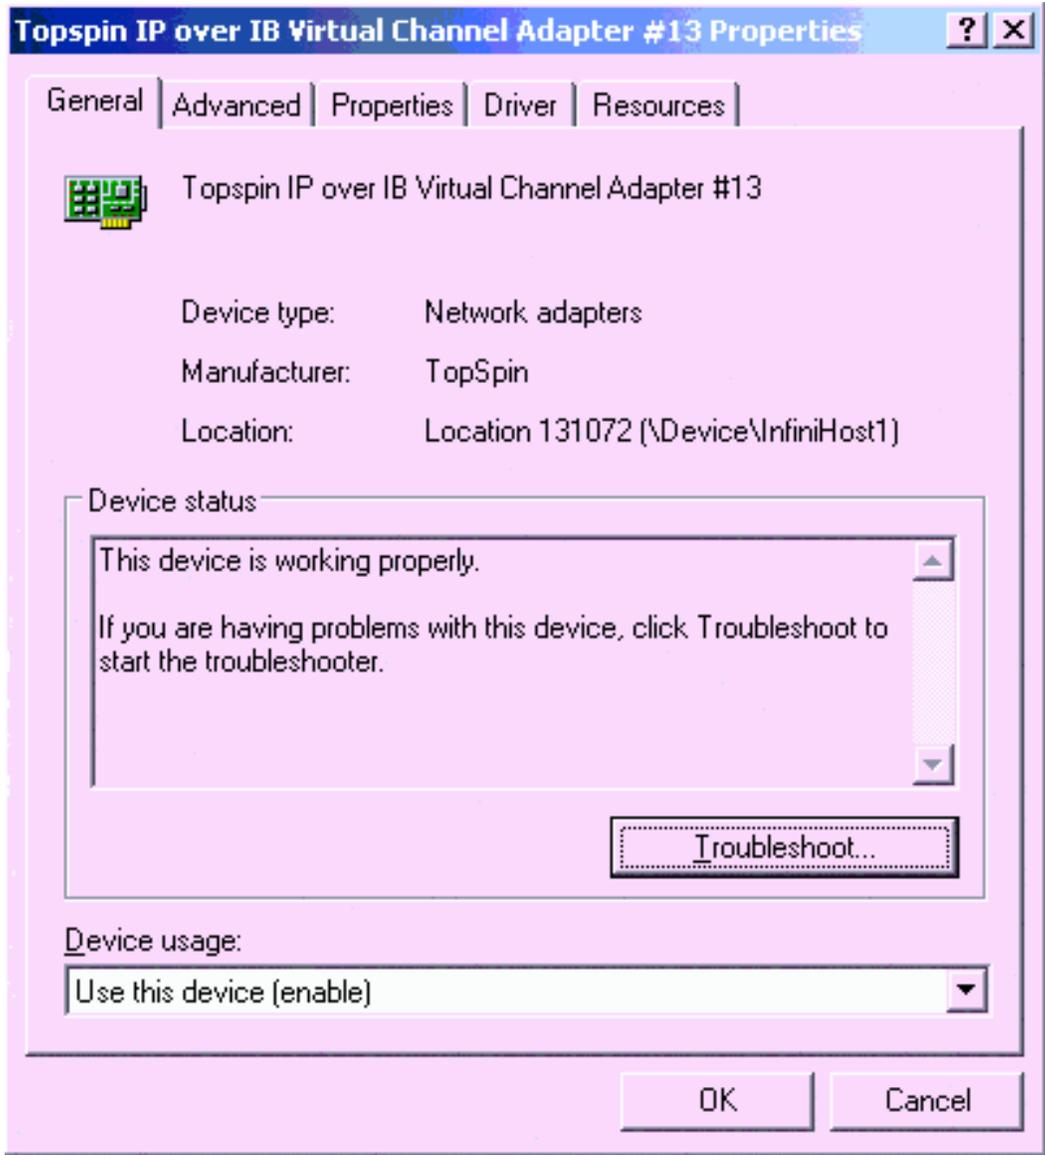
- Expand the **Network Adapters** icon in the right-hand frame of the **Computer Management** window.

Example

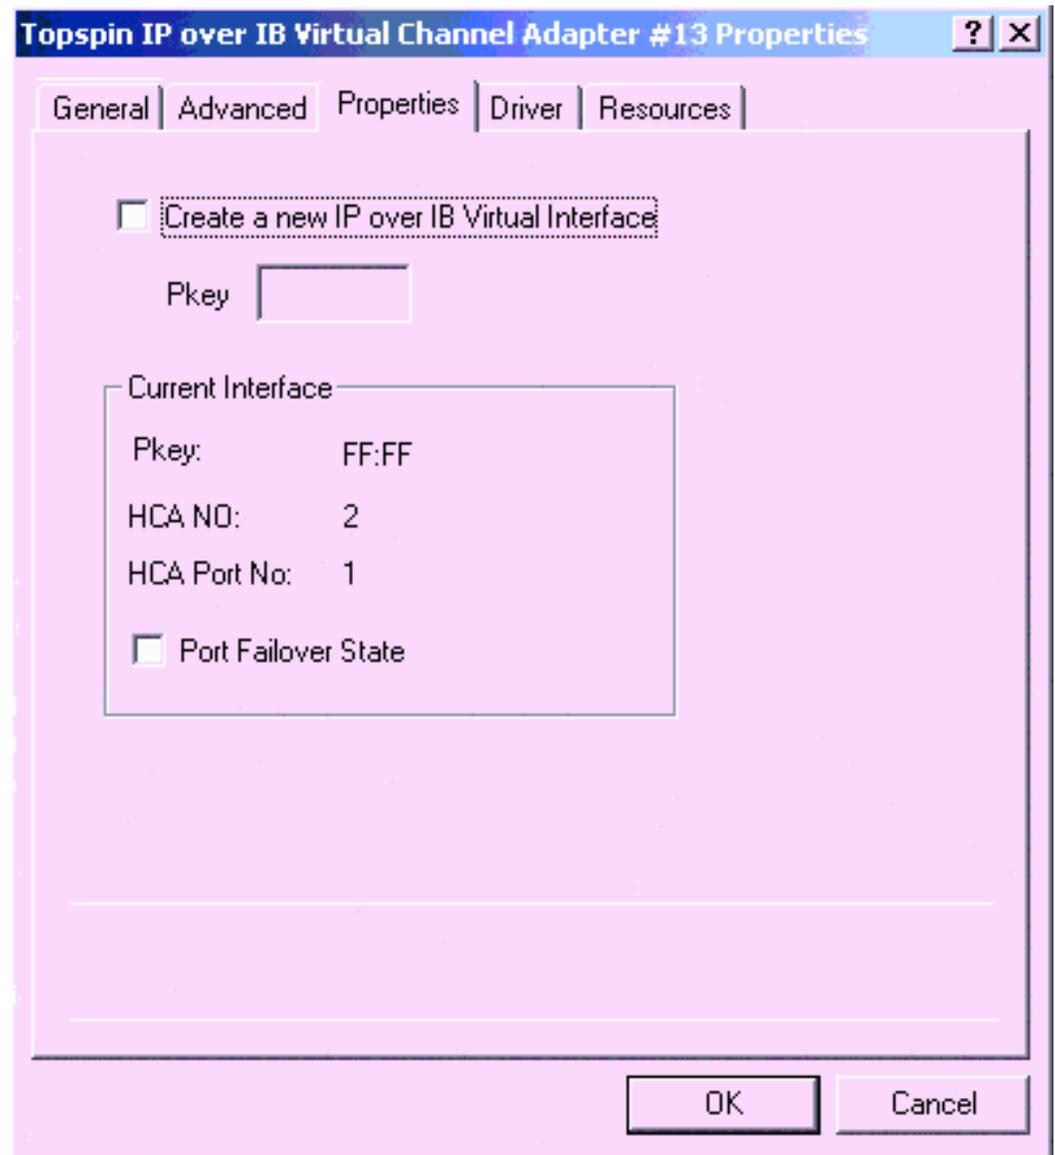


- 5. Right-click a Topspin IP over IB Virtual Adapter (port 0 or port 1) and select **Properties** from the right-click menu. A **Topspin IP over IB Virtual Adapter** properties window appears.

Example

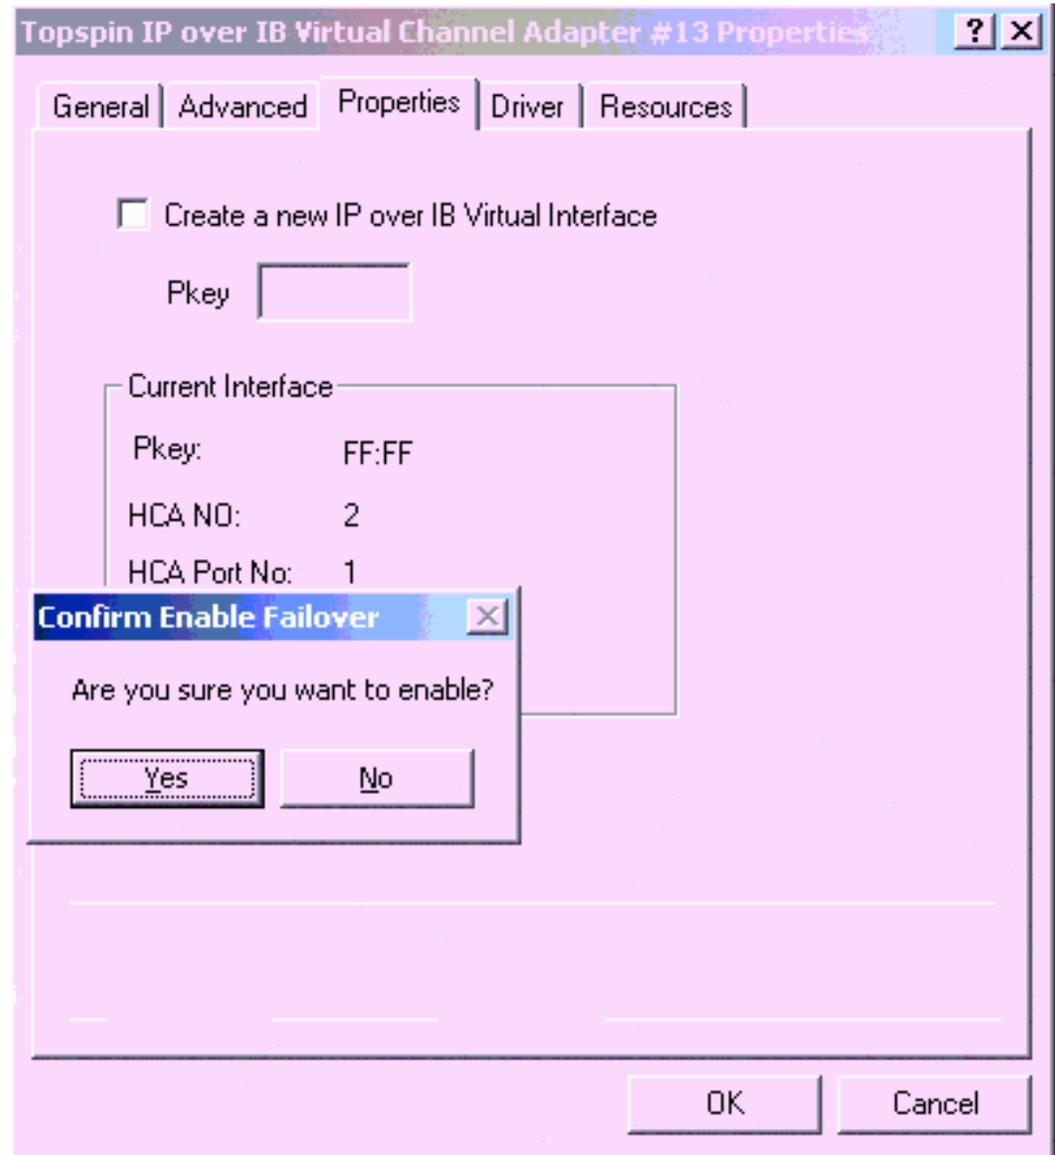


6. Click the **Properties** tab.

Example

7. Check the **Port Failover State** checkbox. The Confirm Enable Failover dialog box appears.

Example



8. Click **Yes** to enable.
9. Click the **Ok** button to close the Properties window or **Cancel** if you do not want to enable port fail-over at this time.

If you click Ok, a dialog box appears to confirm if you want to reboot the host. You must reboot to enable the port-failover capability.



Figure 2-12: Set Port Failover State

10. Click **OK** to reboot host and enable fail-over between ports on the HCA.
11. Log back into the host after it has finished rebooting.
12. Note that only Port 0 appears in the list of IP over IB adapters. Both physical ports are logically configured as a single port.

Name ▲	Type	Status	Device Name	Ow
Make New Connection				
IP over Infiniband Adapter - Port 0	LAN	Enabled	Topspin IP over IB Virtua...	Sys
Local Area Connection	LAN	Enabled	Intel(R) PRO/1000 MT D...	Sys
Local Area Connection 2	LAN	Network cable unplugged	Intel(R) PRO/1000 MT D...	Sys

Figure 2-13: Network Interfaces

Undo HA Ports

To undo HA ports, perform the following steps:

1. Log in to your Windows host.
2. Navigate to your desktop, right-click the **My Computer** icon, and select **Manage** from the right-click menu. The **Computer Management** window opens.
3. Click **Device Manager** in the left-hand frame of the **Computer Management** window.
4. Expand the **Network Adapters** icon in the right-hand frame of the **Computer Management** window.
5. Right-click the Topspin IP over IB Virtual Adapter that you previously merged from 2 physical ports and select **Properties** from the right-click menu. A **Topspin IP over IB Virtual Adapter** properties window appears.
6. Click the **Properties** tab.
7. Uncheck the **Port Failover State** checkbox. The **Confirm Enable Failover** dialog box appears.
8. Click **Yes** to disable HA.
9. Click the **Ok** button to close the Properties window.

If you click OK, a dialog box appears to confirm if you want to reboot the host. You must reboot to disable the port-failover capability.

10. Click **OK** to reboot host and disable fail-over between ports on the HCA.

Configuring Protocols

The following sections appear in this chapter:

- [“Introduction” on page 35](#)
- [“Configure IPoIB” on page 35](#)
- [“Configure SRP” on page 45](#)
- [“Configure MPI \(Linux Only\)” on page 47](#)
- [“Configure SDP \(Linux Only\)” on page 50](#)

Introduction

When you install host drivers, the installation process installs all available drivers on your host. As of this release, the HCA Expansion Card supports all protocols on Linux hosts. The HCA Expansion Card supports IPoIB and SRP on Windows hosts.

Configure IPoIB

Configure IPoIB to run IP traffic over the InfiniBand network. The IPoIB driver starts automatically. Refer to your Linux distribution documentation for additional information about configuring IP addresses.



NOTE: To enable these IPoIB settings across reboots, you must explicitly add these settings to the networking interface startup script. For a sample script, see [“Sample Startup Script” on page 41](#).

Configure IPoIB on Linux

To configure IPoIB on your Linux host, perform the following steps:

1. Log in to your Linux host.
2. Enter the **ifconfig** command with
 - the appropriate IB interface (**ib0** or **ib1**)
 - the IP address that you want to assign to the interface
 - the **netmask** keyword
 - the subnet mask that you want to assign to the interface
 to configure an IB interface.

Example

```
# ifconfig ib0 192.168.0.176 netmask 255.255.252.0
```

3. (Optional) Verify the configuration with the **ifconfig** command with the **ifconfig ib0** argument.

Example

```
1. # ifconfig ib0
2. ib0      Link encap:Ethernet  HWaddr 93:C1:2A:29:33:3E
3.          inet addr:192.168.0.176  Bcast:192.168.0.255  Mask:255.255.252.0
4.          UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
5.          RX packets:2695034 errors:0 dropped:0 overruns:0 frame:0
6.          TX packets:1195933 errors:0 dropped:0 overruns:0 carrier:0
7.          collisions:0 txqueuelen:128
8.          RX bytes:343087140 (327.1 Mb)  TX bytes:67417660 (64.2 Mb)
```

In the preceding example, the IP address of the interface appears as 192.168.0.176 ([line 3](#)).

4. Repeat [step 2](#) and [step 3](#) on the remaining interface.

Example

```
# ifconfig ib1 192.168.0.2 netmask 255.255.255.0
```

Create Subinterfaces

Physical HCA Expansion Card interfaces serve as parent interfaces for logical subinterfaces. The subinterfaces are distinguished by InfiniBand partition assignments. The default Partition Key (p_key), ff:ff, applies to the primary interface.

Subinterfaces are optional. If you do not configure a subinterface, traffic defaults to the parent interface. The partitions (p_keys) provide traffic isolation.

To create subinterfaces, perform the following steps:

1. Create a partition on the Server Switch Module. Alternatively, you can choose to create the partition of the InfiniBand interface first, and then create the partition for the ports on the Server Switch Module. Refer to the *Topspin Element Manager User Guide* for information regarding partitions on the IB switch.
2. Log in to your host.

3. Navigate to the `/usr/local/topspin/sbin` directory.

Example

```
# cd /usr/local/topspin/sbin
```

4. Enter the **ipoibcfg add** command with
 - the parent interface (**ibX**) to which you want to add the subinterface
 - the partition key value that has been created on the Server Switch Module.to create the subinterface.

Example

```
# ./ipoibcfg add ib0 80:02
```

5. Enter the **ifconfig -a** command to verify that you created the interface (line 27 in the following example).

Example

```
1. # ifconfig -a
2. eth0      Link encap:Ethernet  HWaddr 00:09:6B:B5:73:D0
3.           inet addr:10.2.1.176  Bcast:10.2.3.255  Mask:255.255.252.0
4.           UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
5.           RX packets:3969163 errors:0 dropped:0 overruns:0 frame:0
6.           TX packets:17449 errors:0 dropped:0 overruns:0 carrier:0
7.           collisions:0 txqueuelen:1000
8.           RX bytes:435714894 (415.5 Mb)  TX bytes:2257959 (2.1 Mb)
9.           Interrupt:16
10.
11. eth1      Link encap:Ethernet  HWaddr 00:09:6B:B5:73:D1
12.           BROADCAST MULTICAST  MTU:1500  Metric:1
13.           RX packets:0 errors:0 dropped:0 overruns:0 frame:0
14.           TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
15.           collisions:0 txqueuelen:1000
16.           RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
17.           Interrupt:17
18.
19. ib0       Link encap:Ethernet  HWaddr 93:C1:2A:29:33:3E
20.           inet addr:192.168.0.176  Bcast:192.168.0.255  Mask:255.255.252.0
21.           UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
22.           RX packets:2695034 errors:0 dropped:0 overruns:0 frame:0
23.           TX packets:1195933 errors:0 dropped:0 overruns:0 carrier:0
24.           collisions:0 txqueuelen:128
25.           RX bytes:343087140 (327.1 Mb)  TX bytes:67417660 (64.2 Mb)
26.
27. ib0.8002  Link encap:Ethernet  HWaddr 00:00:00:00:00:00
28.           BROADCAST MULTICAST  MTU:2044  Metric:1
29.           RX packets:0 errors:0 dropped:0 overruns:0 frame:0
30.           TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
31.           collisions:0 txqueuelen:128
32.           RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
33.
34. ib1       Link encap:Ethernet  HWaddr DA:C1:2F:17:55:E8
35.           inet addr:192.168.4.176  Bcast:192.168.4.255  Mask:255.255.252.0
36.           UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
37.           RX packets:12 errors:0 dropped:0 overruns:0 frame:0
38.           TX packets:5 errors:0 dropped:0 overruns:0 carrier:0
39.           collisions:0 txqueuelen:128
40.           RX bytes:756 (756.0 b)  TX bytes:384 (384.0 b)
41.
42. lo        Link encap:Local Loopback
43.           inet addr:127.0.0.1  Mask:255.0.0.0
44.           UP LOOPBACK RUNNING  MTU:16436  Metric:1
45.           RX packets:362 errors:0 dropped:0 overruns:0 frame:0
46.           TX packets:362 errors:0 dropped:0 overruns:0 carrier:0
47.           collisions:0 txqueuelen:0
48.           RX bytes:43815 (42.7 Kb)  TX bytes:43815 (42.7 Kb)
```

6. Enter the **cd** command to return to the root directory.

Example

```
# cd
```

7. Configure the new interface just as you would the parent interface (see [“Configure IPoIB on Linux” on page 36](#)).

Example

```
# ifconfig ib0.8002 192.168.5.5 netmask 255.255.252.0
```

Remove a Subinterface

To remove a subinterface, perform the following steps:

1. Bring the subinterface off line. You cannot remove a subinterface until you bring it down.

Example

```
# ifconfig ib0.8002 down
```

2. Navigate to the `/usr/local/topspin/sbin` directory

Example

```
# cd /usr/local/topspin/sbin
```

3. Enter the **ipoibcfg** command with
 - the **del** keyword
 - the primary interface of the subinterface that you want to remove
 - the P_key of the subinterface that you want to remove
 to delete the subinterface.

Example

```
# ./ipoibcfg del ib0 80:02
```

4. (Optional) Enter the **cd** command to return to the root directory

Example

```
# cd
```

5. (Optional) Enter the **ifconfig -a** command to verify that the subinterface no longer appears in the interface list.

Example

```
# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:09:6B:B5:73:D0
          inet addr:10.2.1.176  Bcast:10.2.3.255  Mask:255.255.252.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:4034686 errors:0 dropped:0 overruns:0 frame:0
          TX packets:18863 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:443157905 (422.6 Mb)  TX bytes:2497589 (2.3 Mb)
          Interrupt:16

eth1      Link encap:Ethernet  HWaddr 00:09:6B:B5:73:D1
          BROADCAST MULTICAST  MTU:1500  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:1000
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Interrupt:17

ib0      Link encap:Ethernet  HWaddr 93:C1:2A:29:33:3E
          inet addr:192.168.0.176  Bcast:192.168.0.255  Mask:255.255.252.0
          UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
          RX packets:2695034 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1195933 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:128
          RX bytes:343087140 (327.1 Mb)  TX bytes:67417660 (64.2 Mb)

ib1      Link encap:Ethernet  HWaddr DA:C1:2F:17:55:E8
          inet addr:192.168.4.176  Bcast:192.168.4.255  Mask:255.255.252.0
          UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
          RX packets:12 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:128
          RX bytes:756 (756.0 b)  TX bytes:384 (384.0 b)

lo       Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:378 errors:0 dropped:0 overruns:0 frame:0
          TX packets:378 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:45023 (43.9 Kb)  TX bytes:45023 (43.9 Kb)
```

Sample Startup Script

The file that follows is an example file named `ifcfg-ib0` that resides in `/etc/sysconfig/network-scripts/` on a Linux host.

```
DEVICE=ib0
BOOTPROTO=static
IPADDR=192.168.10.1
NETMASK=255.255.255.0
ONBOOT=yes
```

Configure IPoIB on Windows

1. Click **Start > Settings > Network and Dial-up Connections**. The **Network and Dial-up Connections** window opens (Figure 3-1).

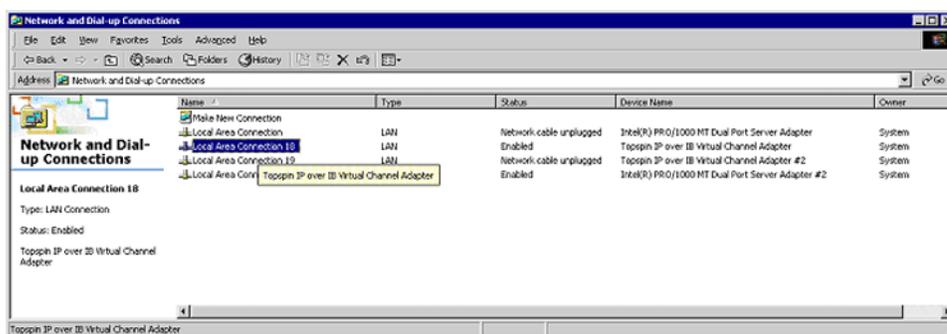
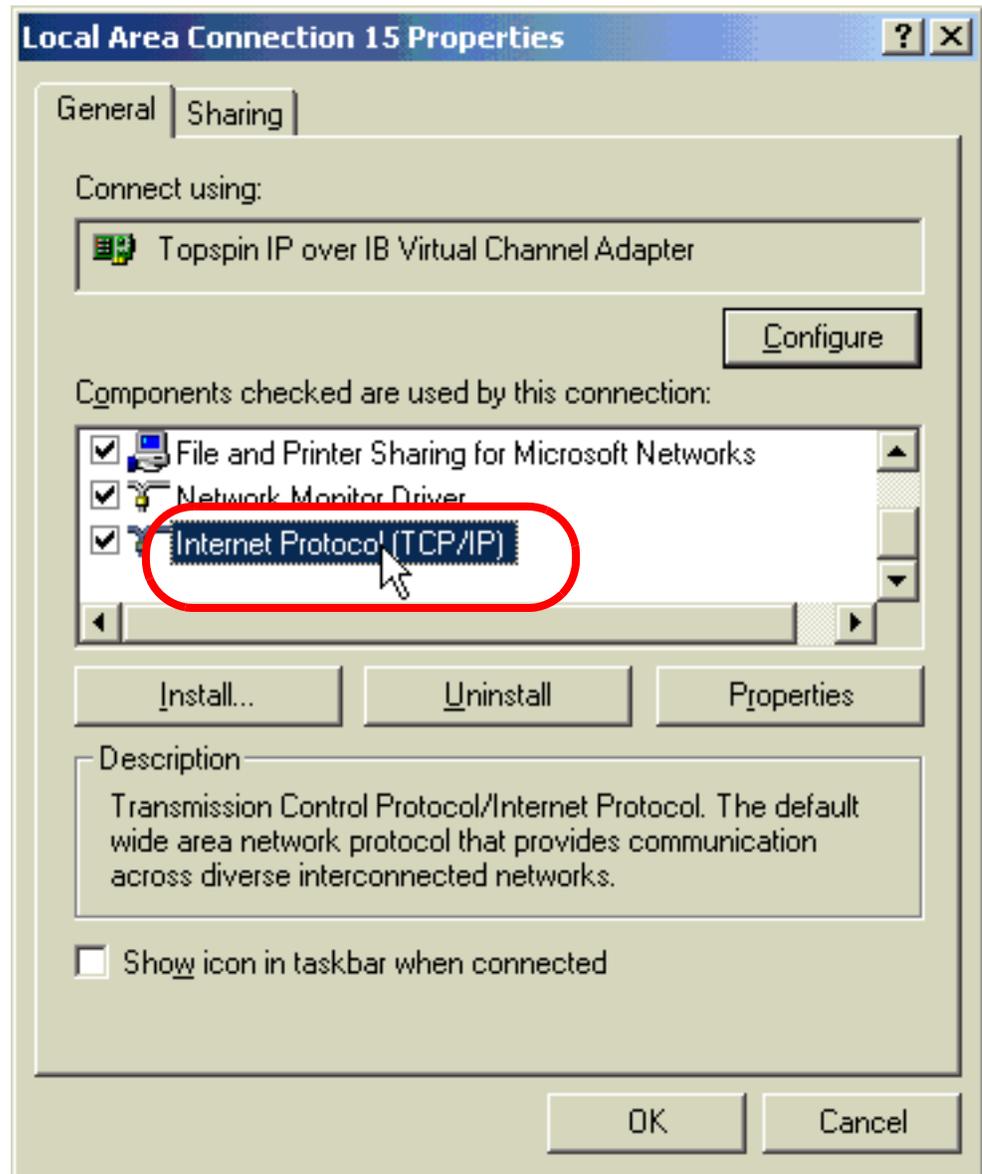


Figure 3-1: Network and Dial-up Connections Window

2. Double-click one of the port entries. (IB port entries display **Topspin IP over IB Virtual Channel Adapter** in the **Device Name** column.) The **Local Area Connection Status** window for the interface opens.
3. Click the **Properties** button. The **Local Area Connection Properties** window opens.

- Click Internet Protocol (TCP/IP) in the **Components checked are used by this connection** field.

Example



- Click the **Properties** button. The **Internet Protocol (TCP/IP) Properties** window opens.
- Click the **Use the following IP address** radio button.



NOTE: If you want your host to obtain an IP address automatically, skip [step 6](#) through [step 9](#)

- Enter an IP address in the **IP address** field
- Enter a mask in the **Subnet mask** field
- Enter a default gateway IP address in the **Default gateway** field.

- Click the **Use the following DNS server addresses** radio button.



NOTE: If you want the host to obtain DNS server addresses automatically, skip [step 10](#) through [step 12](#).

- Enter the IP address of your preferred DNS server in the **Preferred DNS server** field.
- Enter the IP address of your alternate DNS server in the **Alternate DNS server** field.
- Click the **OK** button to return to close the **Internet Protocol (TCP/IP) Properties** window.
- Click the **OK** button to close the **Properties** window.
- Click the **Close** button to close the local area connection window.

Verify the Configuration

- Log in to another host on the IB network.
- Open a command line window and enter the **ping** command with the IP address of the port that you just configured.

Example

```
ping 10.3.0.15
```

- If the hosts ping successfully, the configuration succeeded.

Configure Subinterfaces

To configure an IPoIB subinterface, perform the following steps:

- Log in to your Windows host.
- Right-click the **My Computer** icon on the desktop and click **Manage** in the right-click menu. The **Computer Management** window opens.
- Click the **Device Manager** icon under the **Tree** tab.
- Expand the **Network adapters** branch in the right-hand navigation tree.
- Right-click the Topspin port that you want to serve as the parent interface of the subinterface, then click **Properties** in the right-click menu. The **Topspin IP over IB Virtual Channel Adapter Properties** window opens.
- Click the **Properties** tab.
- Check the **Create a new IP over IB Virtual Interface** checkbox.
- Enter the partition key (P_key) of the subinterface in the **Pkey** field, then click the **OK** button.

Assign an IP Address to a Subinterface

- Create a subinterface, if you have not already done so. Refer to [“Configure Subinterfaces” on page 43](#).

- Click **Start > Settings > Control Panel > Network and Dial-up Connections** to open the **Network and Dial-up Connections** window. The subinterface appears in the list. The partition key of the subinterface appears appended to the device name (see [Figure 3-2](#)).

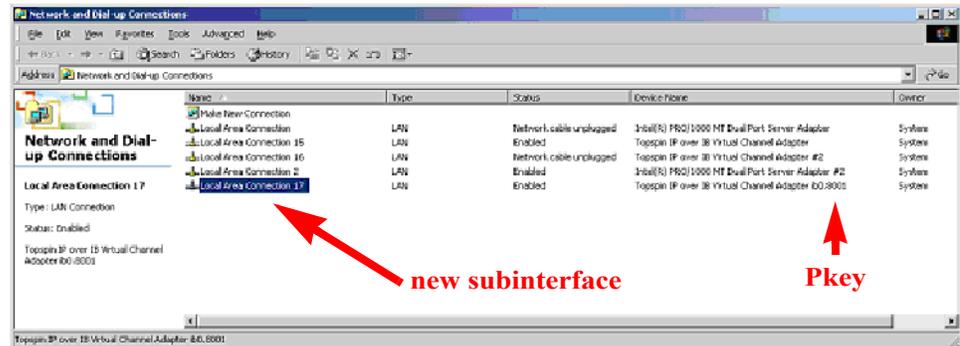
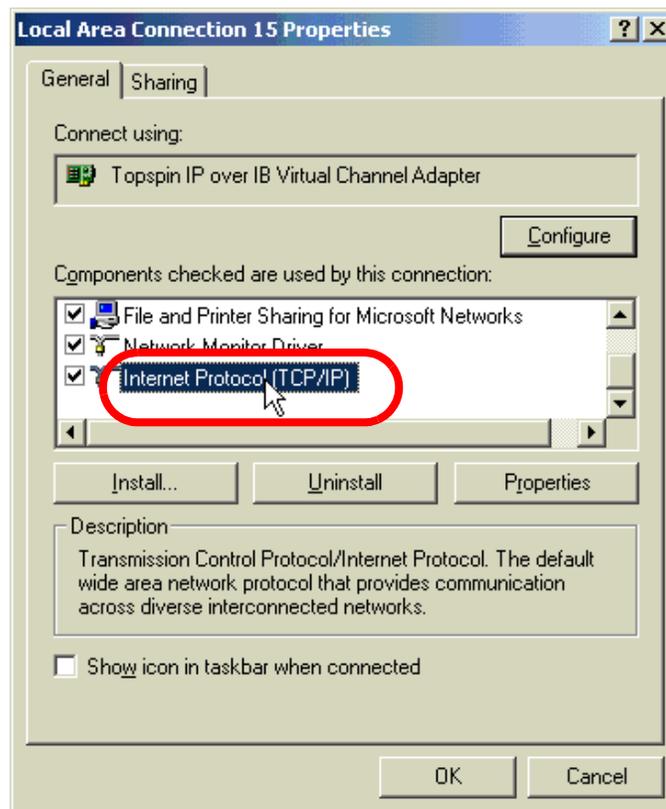


Figure 3-2: New Subinterface

- Double-click the new subinterface. A status window opens.
- Click the **Properties** button. A properties window opens.
- Double-click **Internet Protocol (TCP/IP)**. The **Internet Protocol (TCP/IP) Properties** window opens.

Example



- Assign a static IP address or use DHCP to obtain a dynamic address. The process of assigning or obtaining an IP address for a Partition is the same as configuring an address for a standard interface. See [“Configure IPoIB on Windows” on page 41](#) for instructions.

Remove Subinterfaces

To delete an IPoIB subinterface, perform the following steps:

1. Log in to your Windows host.
2. Right-click the **My Computer** icon on the desktop and click **Manage** in the right-click menu. The **Computer Management** window opens.
3. Click the **Device Manager** icon under the **Tree** tab.
4. Expand the **Network adapters** branch in the right-hand navigation tree.
5. Right-click the subinterface that you want to remove and, in the right-click menu, click **Uninstall**.

Remove Subinterfaces with the Host Utilities CLI

To delete a subinterface key, perform the following steps:

1. Click **Start > Programs > Topspin InfiniBand SDK > Utilities** to open a DOS window.
2. Enter the **ipoibcfg** command with
 - the **del** keyword
 - the HCA port number of the partition (enter **0** for the first port or **1** for the second port)
 - the partition key of the partition.

Example

```
C:\topspin\util>ipoibcfg del 0 80:01
Calling tsDeletePkey
Topspin IP over IB Virtual Channel Adapter ib0.8001

C:\topspin\util>ipoibcfg
Usage : ipoibcfg add <port> <P_Key>
Usage : ipoibcfg del <port> <P_Key>

P_Key should be in the format HH:HH, where H is a hex digit
Example: ipoibcfg add 0 80:0b

C:\topspin\util>
```

Figure 3-3: Deleting a Port From a Partition

Configure SRP

SRP runs on Windows and Linux platforms. To run SRP, simply install the appropriate drivers (see [“Prepare to Install Drivers” on page 17](#)).

To configure your Blade Server to communicate with a storage area network (SAN), you must connect your Server Switch Module to a Topspin Server Switch that includes a Fibre Channel (FC) gateway. To configure your host to access FC storage, refer to the *Fibre Channel Gateway User Guide* that arrives with your FC gateway hardware.

View SRP Devices on Linux

After you configure your host as a SRP host on the appropriate FC-capable Server Switch, use the `cat` command to display attached FC devices.

Example

```
# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: SEAGATE  Model: ST336706LC          Rev: 010A
  Type:   Direct-Access                        ANSI SCSI revision: 03
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: SEAGATE  Model: ST336706LC          Rev: 010A
  Type:   Direct-Access                        ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: DGC      Model: RAID 1              Rev: 0099
  Type:   Direct-Access                        ANSI SCSI revision: 04
Host: scsi2 Channel: 00 Id: 00 Lun: 01
  Vendor: DGC      Model: RAID 5              Rev: 0099
  Type:   Direct-Access                        ANSI SCSI revision: 04
Host: scsi2 Channel: 00 Id: 00 Lun: 02
  Vendor: DGC      Model: RAID 5              Rev: 0099
  Type:   Direct-Access                        ANSI SCSI revision: 04
```

Enable or Disable SRP on Linux

To enable SRP, enter the `ts_srp start` command.

Example

```
# /etc/init.d/ts_srp start
```

To disable SRP, enter the `ts_srp stop` command.

Example

```
# /etc/init.d/ts_srp stop
```

View SRP Devices on Windows

To view SRP devices, perform the following steps:

1. Log in to your Windows host.
2. Navigate to your desktop, right-click the My Computer icon, and select Manage from the right-click menu. The Computer Management window opens.
3. Click Device Manager in the left-hand frame of the window.
4. Expand the Disk drives icon in the Tree tab.
5. Click the Disk Management folder to display the exact layout of storage disks and partitions. The Write signature and Upgrade Disk Wizard opens.
6. Click Cancel to write a signature but bypass Dynamic Disk.



NOTE: If disks do not appear, right-click Device Management and select Rescan Disks.

Configure MPI (Linux Only)



NOTE: MPI software does not appear on your installation CD. You must download the software from the support Web site at <http://support.topspin.com>.

MPI is only available on Linux platforms as of this release. Before you can run MPI, you must establish a SSH connection between two hosts so that you can run commands between the nodes without a log-in or password.

Configure SSH

To configure SSH between two hosts so that a connection does not require a password, perform the following steps:

1. Log in to the host that you want to configure as the local host (hereafter, “host 1”). (The second host serves as the remote host.)

Example

```
login: username
Password: password
Last login: Tue Aug 31 14:52:42 from 10.10.253.115
You have new mail.
[root@qa-bc1-blade4 root]#
```

2. Enter the **ssh-keygen -t rsa** command to generate a public/private RSA key pair. The CLI prompts you for a folder in which to store the key.

Example

```
qa-bc1-blade4:~ # ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
```

3. Press the **Enter** key to store the key in the default directory (/root/.ssh). The CLI prompts you to enter a password.



NOTE: Do not enter a password!

Example

```
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
```

4. Press the **Return** key to bypass the password option. The CLI prompts you to re-enter the password.

Example

```
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
```

5. Press the **Return** key again (once again, omit a password). The CLI displays the fingerprint of the host.

Example

```
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
0b:3e:27:86:0d:17:a6:cb:45:94:fb:f6:ff:ca:a2:00 root@qa-bc1-blade4
qa-bc1-blade4:~ #
```

6. Move to the `.ssh` directory that you created.

Example

```
qa-bc1-blade4:~ # cd .ssh
qa-bc1-blade4:~/.ssh #
```

7. Copy the public key to a file.

Example

```
qa-bc1-blade4:~/.ssh # cp id_rsa.pub authorized_keys
```

8. Log in to the host that you want to configure as the remote host (hereafter “host 2”).

Example

```
login: username
Password: password
Last login: Tue Aug 31 14:52:42 from 10.10.253.115
You have new mail.
[root@qa-bc1-blade5 root]#
```

9. Create a `.ssh` directory in the root directory in host 2.

Example

```
qa-bc1-blade5:~ # mkdir .ssh
```

10. Return to host 1 and copy the file from [step 7](#) to the directory that you created in [step 9](#).

Example

```
qa-bc1-blade4:~/.ssh # scp authorized_keys qa-bc1-blade5:/root/.ssh
```

11. Test your ssh connection.

Example

```
[root@qa-bc1-blade4 root]# ssh qa-bc1-blade5
Last login: Tue Aug 31 14:53:09 2004 from host

[root@qa-bc1-blade5 root]#
```

Edit PATH Variable

1. Establish rsh or ssh connections between two nodes so that you can run commands between a local and remote node without a log-in or password (see “Configure SSH” on page 47).
2. Verify that you do not need to add the compiler to the PATH
3. Add, if required, the following paths to your environment PATH:
 - `/usr/local/topspin/mpi/mpich/bin`
 - `/usr/local/topspin/bin`



NOTE: Optionally, you can add the paths for all users by adding `export PATH=$PATH:/usr/local/topspin/mpi/mpich/bin:/usr/local/topspin/bin` to your `/etc/profile.d` script.

4. Verify that your compiler and MPI script match. Compilers reside in the `/usr/local/topspin/mpi/mpich/bin` directory. GNU compilers use `mpicc` and `mpif77` scripts. Intel compilers use `mpicc.i` and `mpif90.i` scripts.

Perform Bandwidth Test

Before you perform the bandwidth test, configure rsh or ssh on your hosts. To perform the test, perform the following steps:

1. Log in to your local host.
2. Create a **hostfile** text file in the **tmp** directory (`/tmp/hostfile`) that contains the IP addresses of each `ib0` port of the two hosts that will participate in the test.

Example

```
192.168.0.175
192.168.0.176
```

3. Enter the `mpirun_ssh` (or `mpirun_rsh`) command with
 - the **-np** keyword to specify the number of processes
 - the number of processes (integer)
 - the **-hostfile** `/tmp/hostfile` argument
 - the `mpi_bandwidth` command
 - the number of times to transfer the data (integer)
 - the number of bytes to transfer (integer)

to perform the bandwidth test.

Example

```
[root@qa-bc1-blade2 root]# /usr/local/topspin/mpi/mpich/bin/mpirun_ssh -np 2
-hostfile /tmp/hostfile /usr/local/topspin/mpi/mpich/bin/mpi_bandwidth 1000
262144
The authenticity of host 'qa-bc1-blade2 (X.X.X.X)' can't be established.
RSA key fingerprint is 0b:57:f2:c9:dc:cb:ef:67:1c:51:3b:bf:58:8a:35:04.
Are you sure you want to continue connecting (yes/no)?
```

4. Enter **yes** at the prompt to connect to your remote host.

Example

```
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'qa-bc1-blade2,10.2.1.176' (RSA) to the list of known
hosts.
262144 241.250722
[root@qa-bc1-blade2 root]#
```

The output (in the example, **241.250722**) represents available bandwidth, in MB/sec.

Perform Latency Test

Before you perform the bandwidth test, configure rsh or ssh on your hosts. To perform the test, perform the following steps:

1. Log in to your local host.
2. Enter the `mpirun_ssh` command with
 - the **-np** keyword to specify the number of processes
 - the number of processes (integer)
 - the host name of the local host
 - the host name of the remote host
 - the **mpi_latency** command
 - the number of times to transfer the data (integer)
 - the number of bytes to transfer (integer)

to run the latency test.

Example

```
[root@qa-bc1-blade2 root]# /usr/local/topspin/mpi/mpich/bin/mpirun_ssh -np 2 qa-
bc1-blade2 qa-bc1-blade3 /usr/local/topspin/mpi/mpich/bin/mpi_latency 10000 1
1 6.684000
```

The output (in the example, **6.684000**) represents the latency, in microseconds.

Configure SDP (Linux Only)

SDP is only available on Linux platforms as of this release.

Prepare

To prepare, perform the following tasks:

- [“Configure IPoIB” on page 50](#)
- [“Load the Environment Variable” on page 50](#)

Configure IPoIB

Configure IPoIB (see [“Configure IPoIB” on page 35](#)). SDP uses the same IP addresses and interface names as IPoIB.

Load the Environment Variable

Load the installed `libsdp_sys.so` library in one of the following ways:

- Edit the LD_PRELOAD environment variable. Set this to the full path of the library you want to use and it will be pre-loaded.

Example

```
> setenv LD_PRELOAD /lib/libsdp.so
```

- Add the full path of the library into the /etc/ld.so.preload file. The library will be pre-loaded for every executable that is linked with libc.

Example

```
# echo "/lib/libsdp.so" > /etc/ld.so.preload
# more /etc/ld.so.preload
/lib/libsdp.so
```

Convert TCP to SDP in Applications

You can convert TCP to SDP explicitly (manually) or automatically.

Explicit

You can make explicit changes to your programs to replace TCP with SDP. You can choose to

- Replace all TCP instances with SDP.
- Replace specific TCP instances with SDP

To replace all TCP instances with SDP in a C program, add the following lines immediately following the includes:

```
#include "/usr/local/topspin/include/sdp_sock.h"
#undef AF_INET
#define AF_INET AF_INET_SDP
```

To replace specific TCP instances with SDP, perform the following steps:

1. Add the example include that follows to the end of the includes in your program.

```
#include "/usr/local/topspin/include/sdp_sock.h"
```

2. Convert the function calls that you want to use SDP. In the example that follows, [line a](#) uses TCP, but [line b](#) is the same function, converted to use SDP.

Example

- a. `fd = socket(AF_INET, SOCK_STREAM, 0);`
- b. `fd = socket(AF_INET_SDP, SOCK_STREAM, 0);`

Automatic

To configure automatic conversion, configure the ports, IP addresses, or applications that explicitly use SDP by editing the libsdp.conf file.

1. Locate the libsdp.conf file in the /usr/local/topspin/etc directory
2. Make the following modifications:
 - Match on Destination Port

Syntax:

destination ip_addr[/ prefix_length][:start_port [-end_port]]

Example

```
match destination 192.168.1.0/24
```

- Match on Listening Port

Syntax:

listen ip_addr[/prefix_length][:start_port[-end_port]]

Example

```
match listen *:5001
```

- Match on Program Name:

Syntax:

match program program_name*

This uses shell type globs. db2* matches on any program with a name starting with db2. t?p would match on tcp, etc.

Example

```
match program db2*
```

HCA Utilities

The following sections appear in this chapter:

- [“Introduction” on page 53](#)
- [“Run a Self Test” on page 54](#)
- [“Run the vstat Utility to View HCA Attributes” on page 55](#)
- [“Enable or Disable an Interface” on page 57](#)
- [“Run the tvflash -o Utility” on page 58](#)
- [“View and Control an HCA Port on Windows” on page 60](#)
- [“View InfiniBand \(IB\) Statistics” on page 64](#)

Introduction

The sections in this chapter discuss fundamental HCA Expansion Card features. These features address basic usability and provide starting points for troubleshooting.

The HCA Expansion Card includes a self-test that displays information on the health and activity of the HCA Expansion Card. The HCA Expansion Card also includes the vstat utility, which provides a number of different displays to view HCA Expansion Card details. The list of details can include the GUID of the HCA. Lastly, this chapter includes steps that detail how to enable and disable HCA Expansion Card InfiniBand ports.

Run a Self Test

To run an HCA self-test, enter the `hca_self_test` command.

Example

```
[root@qa-bc1-blade2 root]# /usr/local/topspin/sbin/hca_self_test

---- Performing InfiniBand HCA Self Test ----
Number of HCAs Detected ..... 1
PCI Device Check ..... PASS
Host Driver Version ..... rhel3-2.4.21-4.ELsmp-2.0.0-591
Host Driver RPM Check ..... PASS
HCA Type of HCA #0 ..... Cougar
HCA Firmware on HCA #0 ..... v3.01.0000 build 2.0.0.591
HCA Firmware Check on HCA #0 ..... PASS
Host Driver Initialization ..... PASS
Number of HCA Ports Active ..... 2
Port State of Port #0 on HCA #0 ..... UP
Port State of Port #1 on HCA #0 ..... UP
Error Counter Check ..... FAIL
    REASON: found errors in the following counters
        Errors in /proc/topspin/core/cal/port1/counters
            Symbol error counter: 10
Kernel Syslog Check ..... PASS
----- DONE -----
```

[Table 4-1](#) lists and describes the fields in the test output.

Table 4-1: Self Test Fields

Field	Description
Number of HCAs Detected	Number of HCAs on the host that the test recognizes.
PCI Device Check	Confirms that Tavor shows up correctly as a PCI device.
Host Driver Version	Version of the drivers on the host.
Host Driver RPM Check	Confirms that the rpms that are installed are compatible with the host OS.
HCA Type of HCA #0	Displays the “Family” of the HCA (cougar, jaguar, or lion family).
HCA Firmware on HCA #0	Firmware version that runs on the HCA.
HCA Firmware Check on HCA #0	Displays PASS or FAIL.
Host Driver Initialization	Confirms that the IPoIB driver was installed correctly.
Number of HCA Ports Active	Number of enabled ports on the HCA.
Port State of Port #0 on HCA #0	Displays up or down to reflect the status of the port.
Port State of Port #1 on HCA #0	Displays up or down to reflect the status of the port.
Error Counter Check	Displays PASS or FAIL.
Kernel Syslog Check	Displays PASS or FAIL.

Run the vstat Utility to View HCA Attributes

Use the vstat utility to view GUIDs and other HCA attributes. The vstat utility is available on both Windows and Linux platforms.

View Host GUID with vstat Utility on Linux

To view your host GUID on a Linux device, perform the following steps:

1. Log in to your host.
2. Enter the `cd` command to move to the `/usr/local/topspin/bin` directory

Example

```
# cd /usr/local/topspin/bin
```

3. Enter the `vstat --verbose` command. The GUID of the HCA appears in [line 26](#).



NOTE: Abbreviated output appears in the following example.

Example

```
1. [root@qa-bc1-blade2 bin]# ./vstat --verbose
2. 1 HCA found:
3.     hca_id=InfiniHost0
4.     vendor_id=0x02C9
5.     vendor_part_id=0x5A44
6.     hw_ver=0xA1
7.     fw_ver=0x300010000
8.     num_phys_ports=2
9.     max_num_qp=0xfb8      (Maximum Number of QPs supported)
10.    max_qp_ous_wr=0xffff  (Maximum Number of outstanding WR on any WQ)
11.    flags=0x00001476
12.    max_num_sg_ent=0x3b   (Max num of scatter/gather entries for WQE other
13.    than RD)
14.    max_num_sg_ent_rd=0x3b (Max num of scatter/gather entries for RD WQE)
15.    max_num_srq=0x3f0     (Maximum Number of SRQs supported)
16.    max_wqe_per_srq=0xffff (Maximum Number of outstanding WR on any SRQ)
17.    max_srq_sentries=0x3b (Maximum Number of scatter entries for SRQ WQE)
18.    max_num_cq=0xff80    (Max num of supported CQs)
19.    max_num_ent_cq=0x1ffff (Max num of supported entries per CQ)
20.    max_num_mr=0x1ffff0  (Maximum number of memory region supported)
21.    max_mr_size=0xFFFFFFFF (Largest contiguous block of memory regio
22.    n in bytes)
23.    max_pd_num=0x3ffe     (Maximum number of protection domains supported)
24.    page_size_cap=0x1000  (Largest page size supported by this HCA)
25.    max_pkeys=0x40       (Maximum number of partitions supported)
26.    node_guid=00:05:ad:00:00:01:60:44      (Node GUID for this hca)
27.    local_ca_ack_delay=0x10 (Log2 4.096usec Max. RX to ACK or NAK delay)
28.    port_lid=0x0181
```

Run the Basic vstat Utility

To run the vstat utility on Linux, perform the following steps:

1. Log in to your Linux host
2. Enter the `/usr/local/topspin/bin/vstat` command.

Example

```
# /usr/local/topspin/bin/vstat
1 HCA found:
  hca_id=InfiniHost0
  vendor_id=0x02C9
  vendor_part_id=0x5A44
  hw_ver=0xA1
  fw_ver=0x300010000
  num_phys_ports=2
    port=1
    port_state=PORT_ACTIVE
    sm_lid=0x000a
    port_lid=0x0016
    port_lmc=0x00
    max_mtu=2048
    gid_tbl_len=32
    GID[ 0]= fe:80:00:00:00:00:00:00:05:ad:00:00:00:17:18

    port=2
    port_state=PORT_ACTIVE
    sm_lid=0x000a
    port_lid=0x0002
    port_lmc=0x00
    max_mtu=2048
    gid_tbl_len=32
    GID[ 0]= fe:80:00:00:00:00:00:00:05:ad:00:00:00:17:19
```

[Table 4-2](#) lists and describes the fields in the `vstat` command output.

Table 4-2: vstat Fields

Field	Description
<code>hca_id</code>	ID of the HCA.
<code>vendor_id</code>	Vendor ID of the HCA.
<code>vendor_pard_id</code>	ID number of the HCA.
<code>hw_ver</code>	Hardware version of the HCA.
<code>fw_ver</code>	Firmware version that runs on the HCA.
<code>num_phys_ports</code>	Number of physical ports on the HCA.
<code>port</code>	Port on the HCA.  NOTE: In this display, the port appears as 1 or 2, not 0 or 1.
<code>port_state</code>	Displays <code>PORT_ACTIVE</code> or <code>PORT_DOWN</code> to display whether the port is up or down.

Table 4-2: vstat Fields (Continued)

Field	Description
sm_lid	LID of the subnet manager on the switch that connects to the HCA.
port_lid	LID that the subnet manager assigns to the port.
port_lmc	LID mask control (LMC) of the port.
max_mtu	Maximum transmission unit (in bytes) of the port.
gid_tbl_len	GID table length.
GID	Global ID (GID) of the port.

vstat on Windows

To use the vstat utility, perform the following steps:

1. Log in to your Windows host.
2. Click **Start > Programs > Topspin IB SDK > Utilities**. The **Utilities** window opens and displays a list of utilities.
3. Enter the **cd hca** command to move to the hca directory.
4. Enter the **dir** command to view the utilities in the hca directory.
5. Enter the **vstat** command to run the vstat utility. [Table 4-3](#) lists and describes the fields that appear.

Table 4-3: vstat Fields

Field	Description
hca_id	ID of the HCA. As of this release, we only support 1 hca on win2k hosts, so InfiniHost0 appears.
vendor_id	Vendor ID of the HCA.
pard_id	ID number of the HCA.
hw_ver	Hardware version of the HCA.
fw_ver	Firmware version that runs on the HCA.
num_phys_ports	Number of physical ports on the HCA.
port_state	Displays PORT_ACTIVE or PORT_DOWN to display whether the port is up or down.
sm_lid	LID of the subnet manager on the switch that connects to the HCA.
port_lid	LID that the subnet manager assigns to the port.
port_lmc	LID mask control (LMC) of the port.
max_mtu	Maximum transmission unit (in bytes) of the port.

6. Enter the **vstat -v** (verbose) command to view the GUID of the port. The GUID appears in the **node_guid=** field.

Enable or Disable an Interface

When you troubleshoot, verify that your interfaces run successfully. The following sections describe how to bring up and shut down InfiniBand interfaces.

Linux

To bring down an InfiniBand interface on Linux, perform the following steps:

1. Log in to your host.
2. Enter the **ifconfig** command with
 - the port that you want to disable
 - the **down** keyword

Example

```
# ifconfig ib0 down
```

To bring up an InfiniBand interface on Linux, perform the following steps:

1. Log in to your host.
2. Enter the **ifconfig** command with
 - the port that you want to disable
 - the **up** keyword

Example

```
# ifconfig ib0 up
```

Windows

To disable an HCA Expansion Card port on a Windows host, perform the following steps:

1. Click **Start > Settings > Network and Dial-up Connections**. The **Network and Dial-up Connections** window opens.
2. Right-click the interface that you want to bring down.
3. Select **Disable** from the right-click menu.

To enable an HCA Expansion Card port on a Windows host, perform the following steps:

1. Click **Start > Settings > Network and Dial-up Connections**. The **Network and Dial-up Connections** window opens.
2. Right-click the interface that you want to bring up.
3. Select **Enable** from the right-click menu.

Run the tvflash -o Utility

Run the **tvflash -o** command to view boot details. To run this utility, perform the following steps:

1. Enter the **cd** command to move to the **/usr/local/topspin/sbin/** directory.

```
[root@qa-bc1-blade2 root]# cd /usr/local/topspin/sbin/
```

2. Enter the **tvflash -o** command.

```
[root@qa-bc1-blade2 sbin]# ./tvflash -o
auto_upgrade=yes
boot_enable_port_1=yes
boot_enable_port_2=yes
boot_service_scan=yes
boot_type=disable
boot_saved_port=2
boot_saved_ioc_num=1
boot_saved_dgid=fe80:0000:0000:0000:0005:ad00:0000:14aa
boot_saved_service_name=2200:000c:5005:633c
```

Configure Boot Type

Configure boot type to determine the type of remote boot that the HCA executes.

1. Enter the **cd** command to move to the `/usr/local/topspin/sbin/` directory.

```
[root@qa-bc1-blade2 root]# cd /usr/local/topspin/sbin/
```

2. Enter the **tvflash** command with

- the **-o** argument
- the **boot_type=** argument
- the boot type

to configure the boot type on a Linux HCA.

```
[root@qa-bc1-blade2 sbin]# ./tvflash -o boot_type=pxe
Writing [=====]
```

3. (Optional) Enter the **tvflash -o** command to verify the new boot type.

```
[root@qa-bc1-blade2 sbin]# ./tvflash -o
auto_upgrade=yes
boot_enable_port_1=yes
boot_enable_port_2=yes
boot_service_scan=yes
boot_type=pxe
boot_saved_port=2
boot_saved_ioc_num=1
boot_saved_dgid=fe80:0000:0000:0000:0005:ad00:0000:14aa
boot_saved_service_name=2200:000c:5005:633c
```

View and Control an HCA Port on Windows

1. Right-click **My Computer**, then select **Manage** from the right-click menu. The **Computer Management** window opens.
2. Click **Device Manager**, then expand the **InfiniBand** in the right-window.

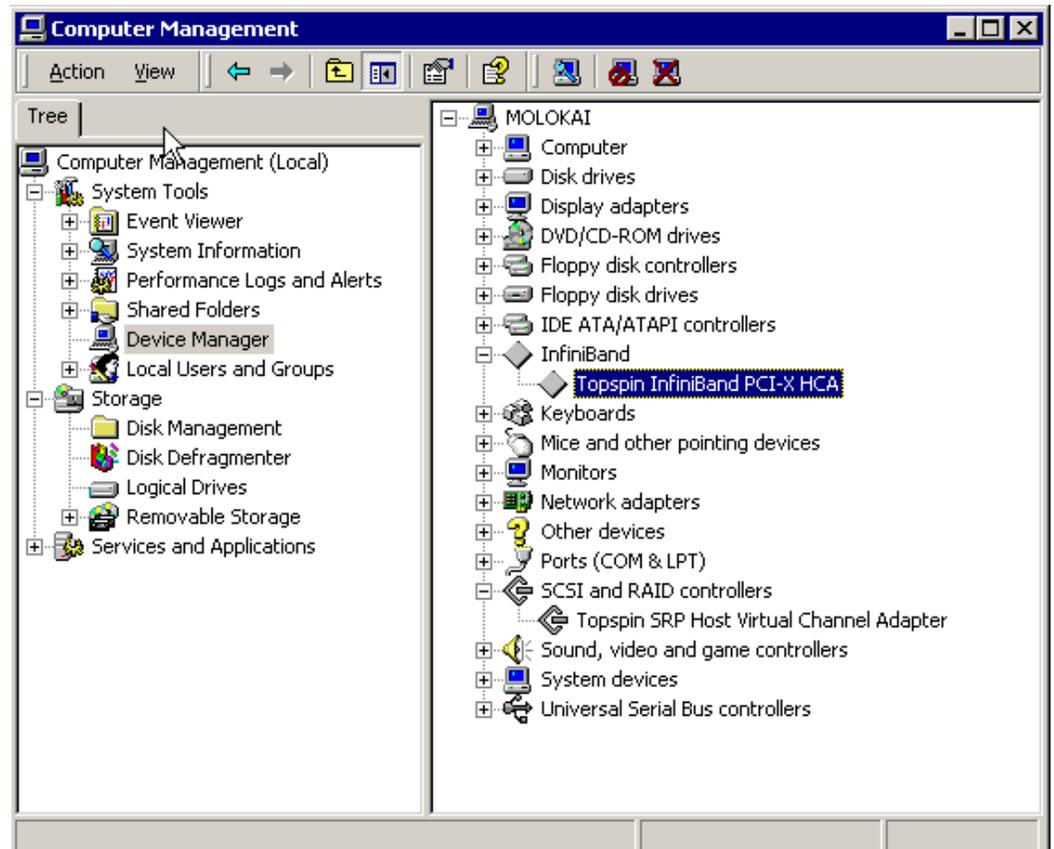


Figure 4-1: Computer Management Window

3. Right-click an InfiniBand HCA entry and select **Properties**. The General tab appears by default.
4. Click the **Diagnostics** tab.

5. Select the port that you want to view or control in the **Port Status and Control for Port #** pulldown menu.

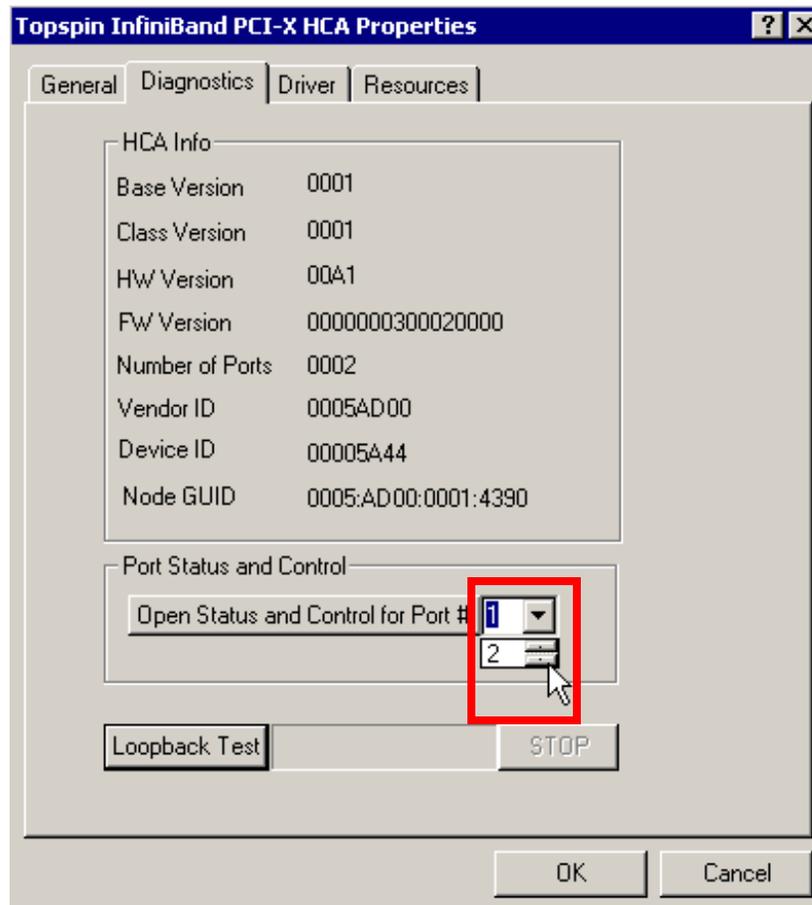


Figure 4-2: Select HCA Port

- Click the **Open Status and Control for Port #** button.

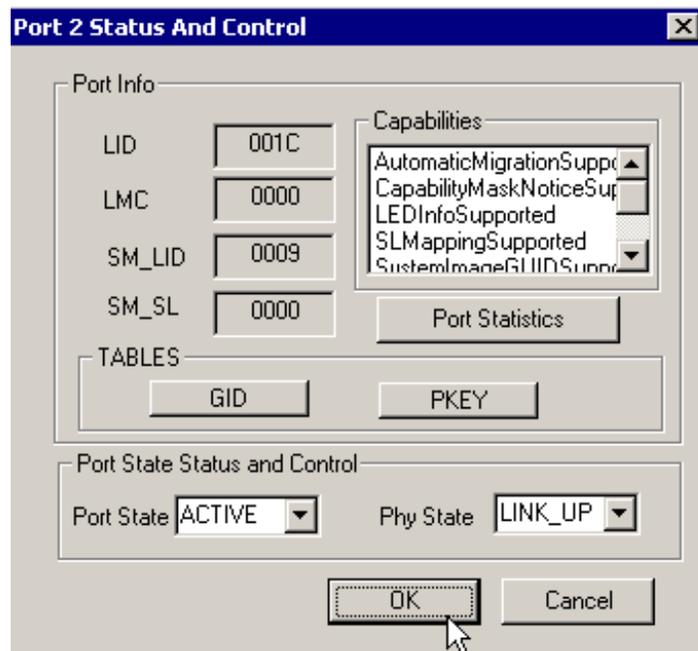


Figure 4-3: Port Status and Control

- View the Logical Identifier (LID) for the HCA port in the **LID** field. The LID is assigned to the port by the active Subnet Manager.
- View the Logical Identifier (LID) for the active Subnet Manager (SM) port in the **SM_LID** field. The LID displayed is that of the Server Switch currently running the active Subnet Manager. In a fabric with multiple Server Switches, only a single switch has an active SM. For example, if a Topspin 360 switch is running the active SM, the LID displayed here would be the LID for the active controller on that switch.

Both ports on the HCA should show the same SM_LID if they are in the same subnet. If the SM_LID is different for each port, it could mean that the ports are in different subnets.
- (Optional) Click the **Port Statistics** button to view detailed statistics for the HCA port.
- Click the **GID** button to view the Global Identifier of the port. Refer to [“View the GID of an HCA Port” on page 62](#).
- Click the **PKEY** button to view the partitions to which the port belongs. The Port Pkey table appears. Refer to [“View the Pkeys to Which a Port Belongs” on page 63](#).
- View the current state of the HCA port in the **Port State** pulldown menu. You can also change the state of the port to ACTIVE or INACTIVE.
- View the current state of the logical link for the port in the **Phy State** field. You can also change the state of the logical link to LINK_UP or DISABLED.
- Click **OK** to save any changes and exit the window.
- Click **Cancel** to exit the window without saving changes.

View the GID of an HCA Port

- Open the **HCA Properties** window. Refer to [“View and Control an HCA Port on Windows” on page 60](#).
- Click the **Diagnostic** tab.
- Click the **GID** button to view the Global Identifier of the port. The GID table appears.

The GID is comprised of the default **subnet prefix** and the **node GUID** of the HCA.

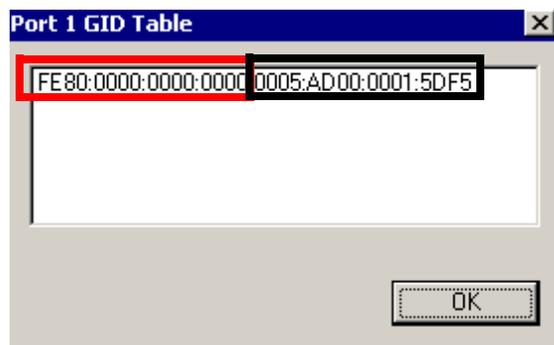


Figure 4-4: Port GID Table

19. Click the **OK** button to exit the GID table.

View the Pkeys to Which a Port Belongs

Refer to the *Element Manager User Guide* for detailed information about Partitions.

20. Open the **HCA Properties** window. Refer to [“View and Control an HCA Port on Windows” on page 60](#).
21. Click the **Diagnostic** tab.
22. Click the **PKEY** button to view the partitions to which the port belongs. The Port Pkey table appears.

The Partition Key (Pkey)s listed are the partitions to which the port has been added.

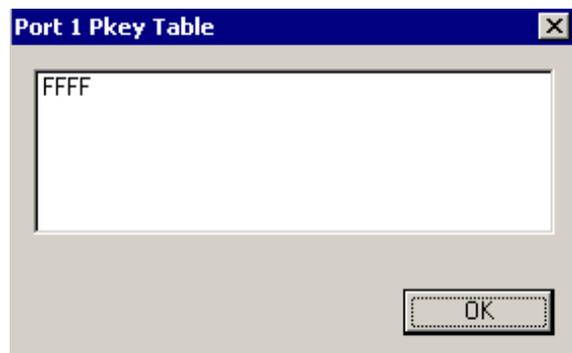


Figure 4-5: Pkey Table

23. Click the **OK** button to exit the Pkey table.

View InfiniBand (IB) Statistics

To display IB Port Statistics:

1. Right-click **My Computer**, then select **Manage** from the right-click menu. The **Computer Management** window opens.
2. Click **Device Manager**, then expand the **InfiniBand** in the right-window.

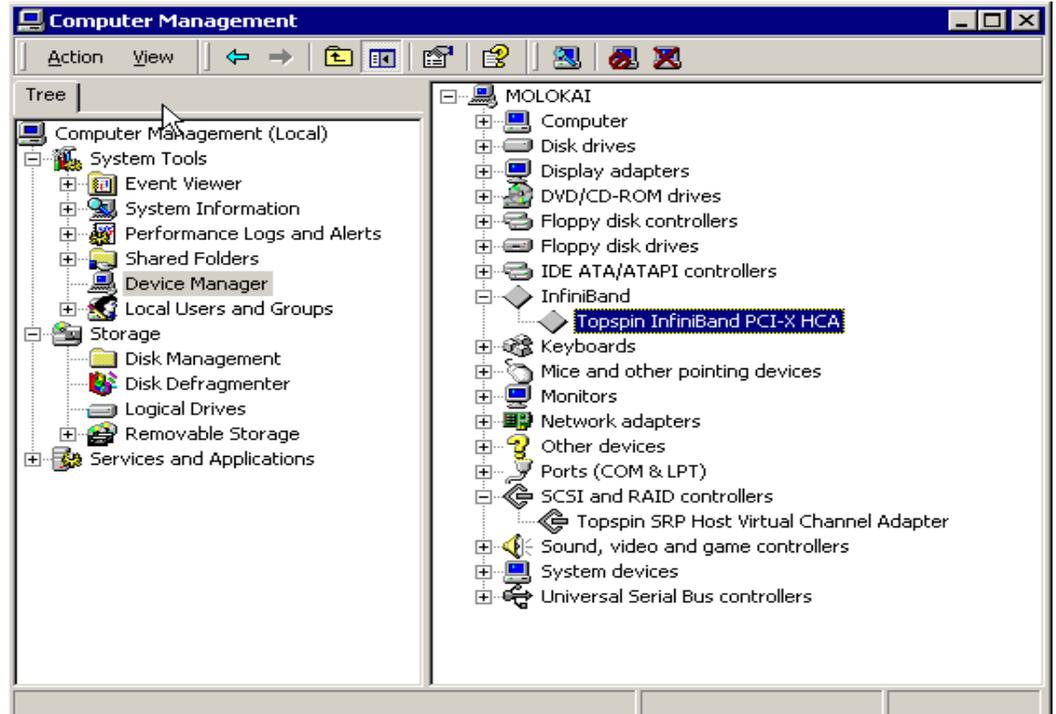


Figure 4-6: Computer Management Window

3. Right-click an InfiniBand HCA entry and select **Properties**. The General tab appears by default.
4. Click the **Diagnostics** tab.

5. Select the port that you want to view or control in the **Port Status and Control for Port #** pulldown menu.

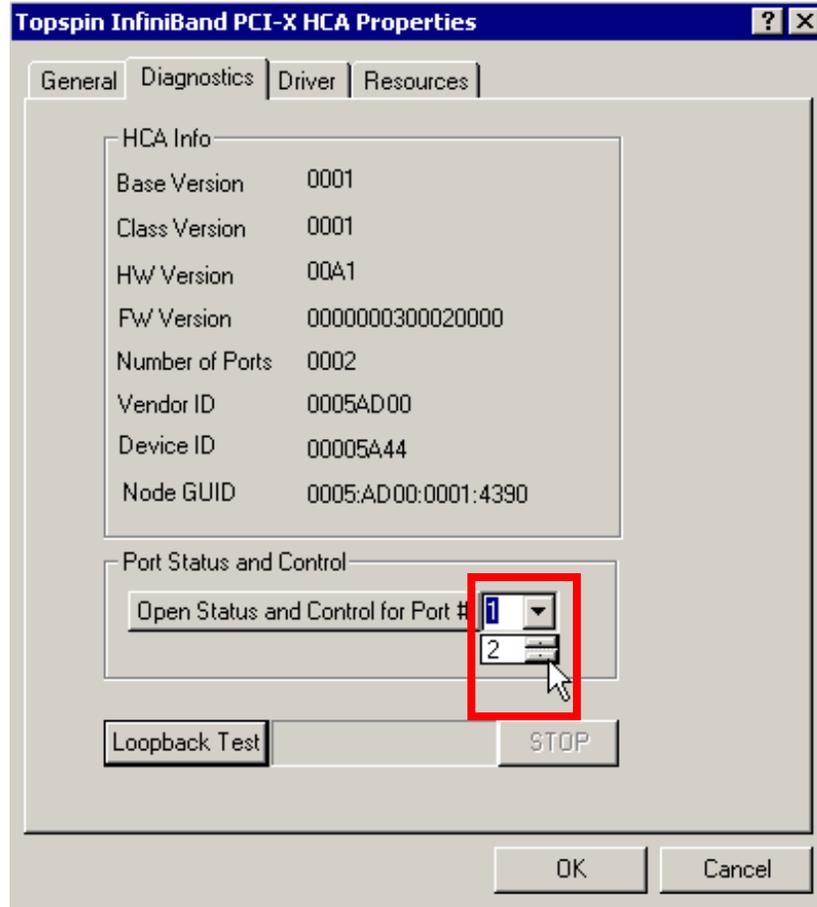


Figure 4-7: Select HCA Port

- Click the **Open Status and Control for Port #** button.

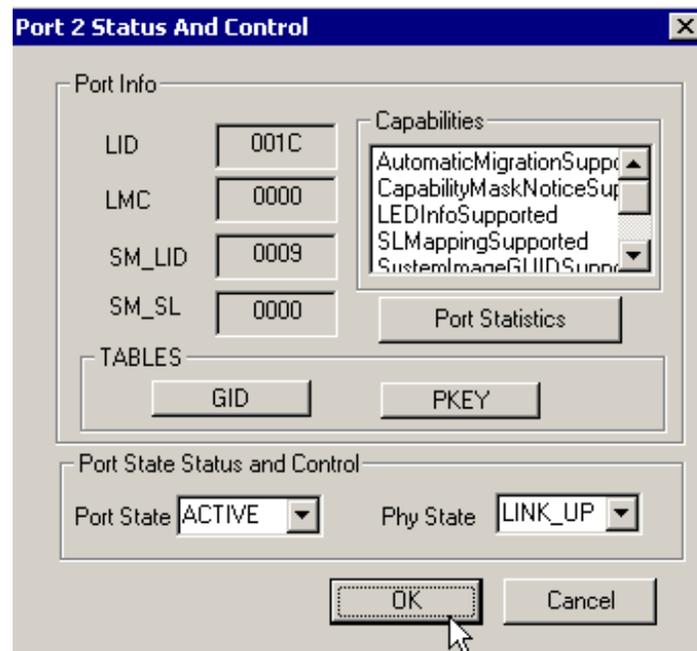


Figure 4-8: Port Status and Control Window

- Click the **Port Statistics** button. The Port Statistics window appears.

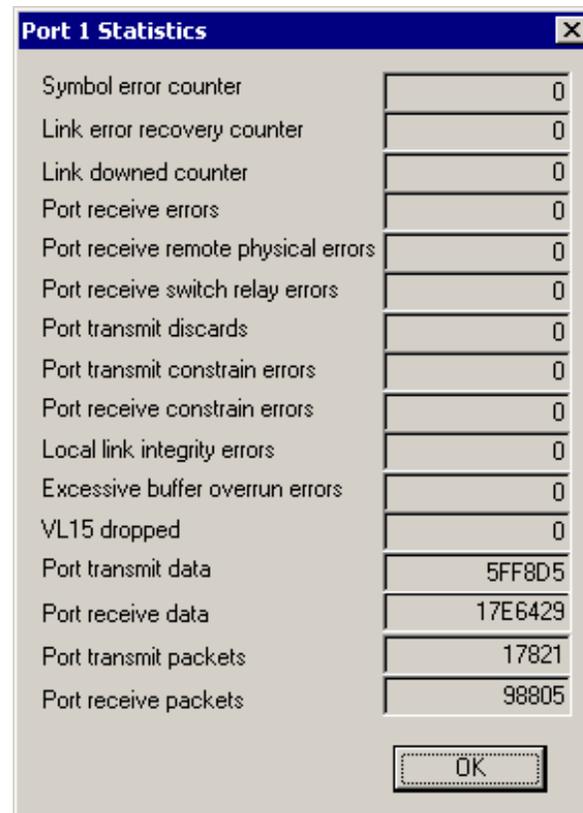


Figure 4-9: Port Statistics

- View the **error** fields. All error-related fields should be 0.

9. View the **Port receive data**, which is displayed in hex.
10. View the **Port transmit packets**, which is displayed in hex.
11. View the **Port receive packets**, which is displayed in hex.

Booting a Host from a Fibre Channel SAN

The following sections appear in this chapter:

- [“Overview” on page 69](#)
- [“Hardware Setup” on page 70](#)
- [“Check HCA Firmware Version” on page 71](#)
- [“BoIB Prerequisites for Windows” on page 71](#)
- [“Configure a Fibre Channel Connection” on page 72](#)
- [“Install a Linux Image onto Fibre Channel Storage” on page 79](#)
- [“Boot an Existing Linux Image” on page 102](#)
- [“Install an Image onto Fibre Channel Storage: Overview” on page 108](#)
- [“Configuration Sample” on page 108](#)

Overview

When hosts boot over an InfiniBand fabric, they boot an image that resides on a remote storage device. Hosts with HCAs that use Boot over IB firmware can boot an image on a Fibre Channel (FC) storage area network (SAN) through Topspin’s Fibre Channel gateway.

To perform a SAN boot with the Boot over IB feature (BoIB-SAN), you must perform the following high-level steps:

1. Install and connect all required hardware. For more, see [“Install Hardware” on page 70](#).
2. Load the correct firmware onto the HCA. You can upgrade Topspin HCA Expansion Cards to run Boot over IB-capable firmware. To upgrade the firmware on your HCA, refer to [“Upgrade Firmware” on page 25](#).



NOTE: Skip this step if your HCA already runs Boot over IB firmware. To determine if your HCA already runs the appropriate software, refer to [“Check HCA Firmware Version” on page 71](#).

3. Configure, on your FC-capable Server Switch, an initiator-target-LUN (ITL) group to connect your Blade Server host to the FC storage. For details, see [“Install a Linux Image onto Fibre Channel Storage” on page 79](#).



NOTE: Skip this step if your environment runs VFrame software. VFrame Director, along with a Topspin Server Switch and gateways, creates a programmable switching fabric for mapping servers to storage and networking resources. For more information, refer to the *Topspin VFrame Director User’s Guide* provided with your VFrame purchase.

4. Load an image onto a Fibre Channel (FC) storage device. For details, see [“Configure a Fibre Channel Connection” on page 72](#).



NOTE: Skip this step if an image already resides on your storage.

5. Boot the remote image. For details, see [“Boot an Existing Linux Image” on page 102](#).

Hardware Setup

A SAN boot with the Boot over IB feature requires the following hardware:

- a Blade Server host with a Topspin HCA Expansion Card
- a Topspin Server Switch Module
- a Topspin Server Switch with a Fibre Channel gateway
- a Fibre Channel storage device (connected, optionally, through a SAN switch or switch fabric)

Install Hardware

To create the physical environment for your SAN boot, perform the following high-level steps:

1. Install the Topspin HCA Expansion Card into your Blade Server host.



NOTE: Only Topspin “Cougar,” “Cougar Cub,” and “Lion Cub” family HCAs support the Boot over IB feature. To verify that your HCA supports BoIB, run the **tvflash -i** command and verify that **Cougar** appears in the output. For complete **tvflash** command syntax, see [“Check HCA Firmware Version” on page 71](#).

2. Install the Topspin Server Switch Module in your BladeCenter unit.
3. Connect the Server Switch Module to an InfiniBand port on your Server Switch.
4. Install the Topspin Fibre Channel gateway into your Server Switch. For details, refer to the *Fibre Channel Gateway User Guide*.
5. Connect the Fibre Channel gateway to a FC storage device or to a SAN fabric with a FC storage device.

Check HCA Firmware Version

This section explains how to check the version of the firmware that runs on your HCA Expansion Card.

Check Version on Linux

To verify that your HCA provides the Boot over IB feature, enter the **tvflash -i** command from the host CLI. If **.Boot** appears in the description of your HCA (see [line 3](#) in the following example), the HCA includes the BoIB feature.



NOTE: Even if your HCA runs 2.0 firmware, it may not run the Boot over IB feature.

If your HCA does not run Boot over IB firmware, update the HCA. For instructions, refer to the *Host Channel Adapter User Guide*.

Example

```
1. [root@test root]# /usr/local/topspin/sbin/tvflash -i
2. HCA #0: Found MT23108, Cougar, revision A1 (firmware autoupgrade)
3.   Primary image is v3.01.0000 build 2.0.0.0, with label 'CA.Cougar.A1..Boot'
4.   Secondary image is v3.01.0000, with label 'HCA.Cougar.A1'
```

Check Version on Windows

To verify that your HCA provides the Boot over IB feature, perform the following steps:

1. Complete the steps under “[Hardware Setup](#)” on [page 70](#).
2. Log in to your Windows 2000 host.
3. Click **Start > Programs > Topspin Infiniband SDK > Utilities**. The **Utilities** window opens.
4. Enter the **tvflash -i** command at the prompt. If **.Boot** appears in the description of your HCA, the HCA includes the BoIB feature. If **.Boot** does not appear, upgrade your firmware.

BoIB Prerequisites for Windows

- Install software on your host (such as Paragon Partition Manager) that lets you copy your local drive to your remote FC storage.
- Install *only one* HCA in your host.
- Configure your SRP host as an initiator on your Server Switch (see *Topspin Fibre Channel Gateway User Guide*).
- (Optional) Configure your Server Switch to grant the host access to only one logical unit (LU) on the FC storage (see *Topspin Fibre Channel Gateway User Guide*). When you restrict access to one LU, you guarantee that you install the image to the correct LU.
- Configure the boot LUN with Element Manager (see *Topspin Fibre Channel Gateway User Guide*). You must configure the switch with the LU that you want the host to boot.
- Size your C drive so that it is the size of your LU or smaller, or verify that your LU is larger than your C drive.

Configure a Fibre Channel Connection

- [“Use Element Manager and a Host with a Remote Boot HCA” on page 72](#)
- [“Use the CLI and a Host with a Remote Boot HCA” on page 73](#)

Use Element Manager and a Host with a Remote Boot HCA

These instructions assume that

- The HCA Expansion Card in your Blade Server host runs Boot over IB feature firmware.
- You have installed all hardware components successfully.

To install an image onto storage, perform the following steps:

1. (Recommended) Restrict global access.



NOTE: If you do not restrict access, only the first seven LUNs that your host can access will appear in the boot menu.



NOTE: After you install an image, you must deny your host access to all LUNs that contain bootable images (except the LUN that stores the image that you want to boot). Hosts might not boot successfully when they can access multiple LUNs that store bootable images. If, after you install the image, you grant your host access to additional LUNs that do not contain bootable images, your host continues to boot successfully.

- a. Launch Element Manager and open your FC-capable Server Switch.
 - b. From the **FibreChannel** menu, select **Storage Manager**. The **Storage Manager** window opens.
 - c. Click the **Storage** folder at the top of the left-hand navigation tree.
 - d. In the **Gateway Port Access** field, check the **Restricted** checkbox.
 - e. In the **LUN Access** field, check the **Restricted** checkbox.
 - f. Click the **Apply** button.
2. Configure an ITL.



NOTE: To install an image with VFrame, do not configure ITLs manually.

- a. Boot your host and press the **x** key when the **Type ‘x’ to configure boot options** prompt appears during the boot process. The GUID of the host appears in the host display. Record the GUID and leave the host at this screen. You will refresh the screen later in this process.
- b. Return to the **Storage Manager** window of Element Manager and click the **SRP Hosts** folder in the navigation tree.
- c. Click the **Define New** button. The **Define New SRP Host** window opens.
- d. Enter the GUID of your host (from [step a](#)) in the **Host GUID** field, or select it from the pulldown menu (if available).
- e. (Optional) Enter a text description for the host in the **Description** field.
- f. Click the **Next >** button.

- g. Click the **Finish** button. The new initiator appears under the **SRP Hosts** folder.
- h. Click the new initiator under the **SRP Hosts** folder, then click the **Targets** tab.
- i. Verify that the WWPN of the target onto which you want to install the image appears in the list.
- j. Double-click the WWPN (from [step i](#)) in the list of the target onto which you want to install the image. The **IT Properties** window opens.
- k. Click the ... button next to the **Port Mask** field. The **Select Ports** window opens.
- l. Click the button of the port that connects the host to the storage, then click **OK**.
- m. Click the **Apply** button in the **IT Properties** window, then close the window.
- n. Click the **LUN Access** tab, then click the **Discover LUNs** button. The targets that your Server Switch sees appear in the **Available LUNs** column.
- o. Click the target (from [step i](#)) onto which you want to install an image, then click the **Add >** button.
- p. Click the **Apply** button.
- q. (Optional) If you plan to install the bootable image to a LU of a RAID device, perform the following steps to define the bootable LU as LUN 0:
 1. In the **Accessible LUNs** field, double-click the LU to which you want to install the image. The **ITL Properties** window opens.
 2. In the **SRP LUN ID** field, enter **00:00:00:00:00:00:00:00**.
 3. Click the **Apply** button, then close the window.
- r. Return to your host and press the **r** key to refresh your view. The WWPN of the target (from [step i](#)) appears in the display, beneath the GUID.
- s. (Optional but recommended) Configure the Well Known Service Name.
 1. Return to the **Storage Manager** window of Element Manager, click your SRP host, then click the **General** tab.
 2. From the **Boot Target WWPN** field, select the boot target from the pulldown menu or manually enter it.
 3. From the **Boot FC LUN** field, select the boot LUN from the pulldown menu or manually enter it.
 4. Click the **Apply** button.

Use the CLI and a Host with a Remote Boot HCA

These instructions assume that

- The HCA in your host runs Boot over IB feature firmware.
- You have installed all hardware components successfully.
- Only one HCA resides in your host.

To install an image onto storage, perform the following steps:

1. (Recommended) Restrict global access.



NOTE: If you do not restrict access, only the first seven LUNs that your host can access will appear in the boot menu.



NOTE: After you install an image, you must deny your host access to all LUNs that contain bootable images (except the LUN that stores the image that you want to boot). Hosts do not boot

successfully when they can access multiple LUNs that store bootable images. If, after you install the image, you grant your host access to additional LUNs that do not contain bootable images, your host continues to boot successfully.

- a. Log into your Server Switch.

Example

```
Login: super
Password: xxxxxx
```

- b. Enter the **enable** command to enter Privileged Exec mode.

Example

```
Topspin-360> enable
Topspin-360#
```

- c. Enter the **configure terminal** command to enter Global Configuration mode.

Example

```
Topspin-360# configure terminal
Topspin-360 (config)#
```

- d. Enter the **fc srp-global** command with
 - the **lun-policy** keyword
 - the **restricted** keyword
 to restrict LUN access to initiators for all new ITs and ITLs.

Example

```
Topspin-360 (config)# fc srp-global lun-policy restricted
```

- e. Enter the **fc srp-global** command with
 - the **gateway-portmask-policy** keyword
 - the **restricted** keyword
 to restrict portmask access to initiators for all new ITs and ITLs.

Example

```
Topspin-360 (config)# fc srp-global gateway-portmask-policy restricted
```

2. Configure an ITL.



NOTE: To install an image with VFrame, do not configure ITLs manually.

- a. Boot your host and press the **x** key when the **Type ‘x’ to configure boot options** prompt appears during the boot process. The GUID of the host appears in the host display. Record the GUID and leave the host at this screen. You will refresh the screen later in this process.
- b. Return to your Server Switch CLI and enter the **fc srp initiator** command with
 - the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00**)
 - the **auto-bind** keyword

to add your host to the configuration file and create WWPNs, on the Fibre Channel gateway, that point to the host.

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c  
00:00:00:00:00:00:00:00 auto-bind
```

- c. (Optional) Enter the **fc srp initiator** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
 - the **description** keyword
 - a text description, enclosed in quotation marks
- to assign an easily-recognizable identifier to the initiator.

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c  
00:00:00:00:00:00:00:00 description "test-initiator"
```

- d. Enter the **exit** command to return to Privileged Exec mode.

Example

```
Topspin-360 (config) # exit  
Topspin-360#
```

- e. Enter the **show fc srp initiator** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)

to view the WWPNs of the host.

Example

```
Topspin-360# show fc srp initiator
```

```
=====
SRP Initiators
=====
      guid: 00:05:ad:00:00:00:17:3c
      extension: 00:00:00:00:00:00:00:00
      description: test-initiator
      wwnn: 20:02:00:05:ad:00:00:00
      credit: 0
      active-ports: none
      pkeys:
      bootup-target: 00:00:00:00:00:00:00:00
      bootup-lu: 00:00:00:00:00:00:00:00
      action: auto-bind
      result: success
      wwpns: port      wwpn      fc-addr
              2/1      20:02:00:05:ad:20:00:00  00:00:00
              2/2      20:02:00:05:ad:24:00:00  00:00:00
              3/1      20:02:00:05:ad:30:00:00  00:00:00
              3/2      20:02:00:05:ad:34:00:00  00:00:00
              4/1      20:02:00:05:ad:40:00:00  00:00:00
              4/2      20:02:00:05:ad:44:00:00  00:00:00
              5/1      20:02:00:05:ad:50:00:00  00:00:00
              5/2      20:02:00:05:ad:54:00:00  00:00:00
              6/1      20:02:00:05:ad:60:00:00  00:00:1d
              6/2      20:02:00:05:ad:64:00:00  00:00:1d
              7/1      20:02:00:05:ad:70:00:00  00:00:00
              7/2      20:02:00:05:ad:74:00:00  00:00:00
              8/1      20:02:00:05:ad:80:00:00  00:00:00
              8/2      20:02:00:05:ad:84:00:00  00:00:00
              9/1      20:02:00:05:ad:90:00:00  00:00:00
              9/2      20:02:00:05:ad:94:00:00  00:00:00
              10/1     20:02:00:05:ad:a0:00:00  00:00:00
              10/2     20:02:00:05:ad:a4:00:00  00:00:00
              11/1     20:02:00:05:ad:b0:00:00  00:00:00
              11/2     20:02:00:05:ad:b4:00:00  00:00:00
              12/1     20:02:00:05:ad:c0:00:00  00:00:00
              12/2     20:02:00:05:ad:c4:00:00  00:00:00
              13/1     20:02:00:05:ad:d0:00:00  00:00:00
              13/2     20:02:00:05:ad:d4:00:00  00:00:00
```

- f. Record the WWPN of the port that connects the host to the storage.
- g. Enter the **show fc srp initiator-wwpn-view** command with
 - the WWPN of the port (from [step f](#))
 - the **target** keyword

to verify that the host can see the target. If the WWPN of the target does not appear in the command output, give the gateway some time to configure the ITs and retry the command.

Example

```
Topspin-360# show fc srp initiator-wwpn-view 20:02:00:05:ad:60:00:00 target
```

```
=====
SRP Targets Accessible to Initiator Via Port WWN 20:02:00:05:ad:60:00:00
=====
```

```

          wwpn: 21:00:00:04:cf:f6:c2:ab
          wwnn: 20:00:00:04:cf:f6:c2:ab
description: SRP.T10:21000004CFF6C2AB
          ioc-guid: 00:05:ad:00:00:00:15:1a
service-name: SRP.T10:21000004CFF6C2AB
protocol-ids: 04:00:00:00:00:00:00:00:00:00
          fc-address: 00:00:23
          mtu: 0
connection-type: nl-port
physical-access: 6/1
```

- h. Enter the **configure terminal** command to return to Global Configuration mode.

Example

```
Topspin-360# configure terminal
Topspin-360 (config) #
```

- i. Enter the **no fc srp it** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00**)
 - the WWPN of the target onto which you want to install the image
 - the **gateway-portmask-policy** keyword
 - the **restricted** keyword
 - the FC gateway port that connects the host to the storage (in slot#/port# format) to grant the host access to one target—the target from which your host will boot—through the port.

Example

```
Topspin-360 (config) # no fc srp it 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00:00:00 21:00:00:04:cf:f6:c2:ab gateway-portmask-policy
restricted 6/1
```

- j. Enter the **fc srp initiator** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00**)
 - the **discover-itl** keyword
- to configure ITLs to which the host belongs. The ITLs inherit the portmask policy of the IT (from [step i](#)).

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00:00:00 discover-itl
```

- k. Enter the **no fc srp itl** command with
- the GUID of the host (from [step a](#))

Install a Linux Image onto Fibre Channel Storage

This section explains how to install a bootable image on Fibre Channel storage with the following resources:

- [“Install and Boot Red Hat 9, Red Hat Enterprise Linux 3.0 or Red Hat Enterprise Linux 4.0” on page 79](#)
- [“Install and Boot Red Hat Enterprise Linux 2.1” on page 86](#)
- [“Install and Boot SUSE Linux Enterprise Server 8 Service Pack 3” on page 88](#)
- [“Install and Boot SUSE Linux Enterprise Server 9 Service Pack 1” on page 95](#)



NOTE: To ensure a successful boot, install with a separate /boot partition during the standard Linux setup.

Install and Boot Red Hat 9, Red Hat Enterprise Linux 3.0 or Red Hat Enterprise Linux 4.0

You can now use a Kickstart configuration file with the `driverdisk` parameter in it to access the BoIB CD over the network.

Example

```
driverdisk --source=nfs:10.0.0.35:/qadata/home/releng/builds/Linux-3.0.0/build173/image/topspin-boib-3.0.0-173.iso
```

You must use an ISO image, not a directory.



NOTE: There is a long delay while the ISO image is mounted and accessed.



NOTE: On RHEL3 and RHEL4, if you install over NFS, you need to use Kickstart with `driverdisk`, or put our ISO image in same directory as Red Hat ISO images.

The following example begins the installation from a CD, but completes the installation over NFS.

1. Place the first installation CD in your host and boot the server from the CD.
2. (RedHat only) At your host CLI, exit interactive mode.
3. (RedHat only) At the **boot** prompt, enter the **linux dd** or **linux dd askmethod** command to prompt the installer to run the Topspin driver CD.

Enter **linux dd askmethod** if you want to do the rest of the installation over NFS.

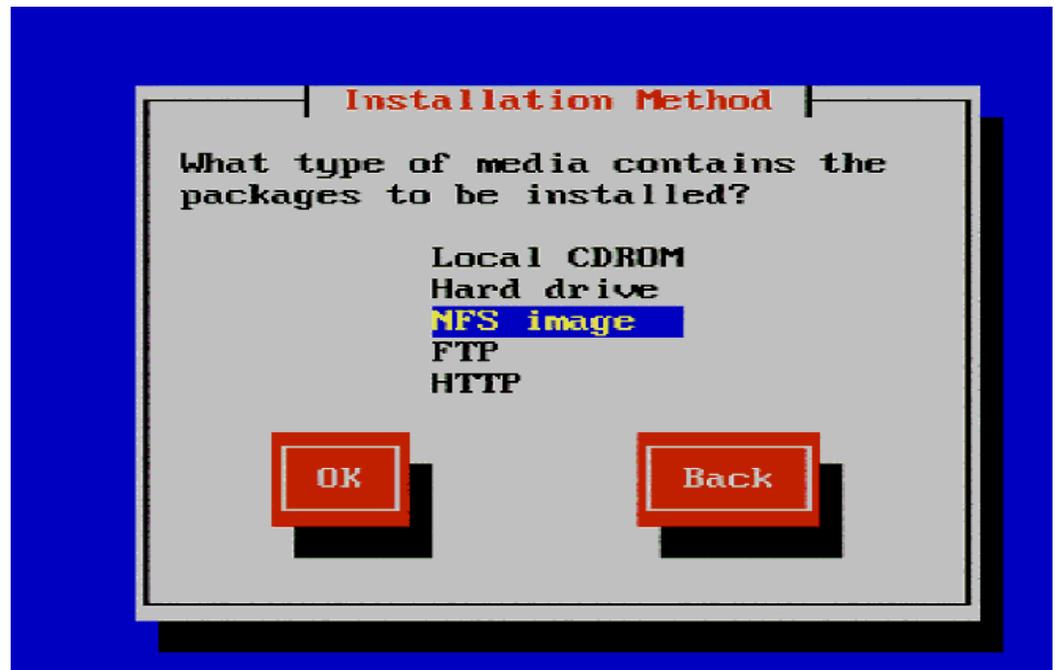


Figure 5-1: Begin the Linux and Red Hat Installation

The installation begins loading.

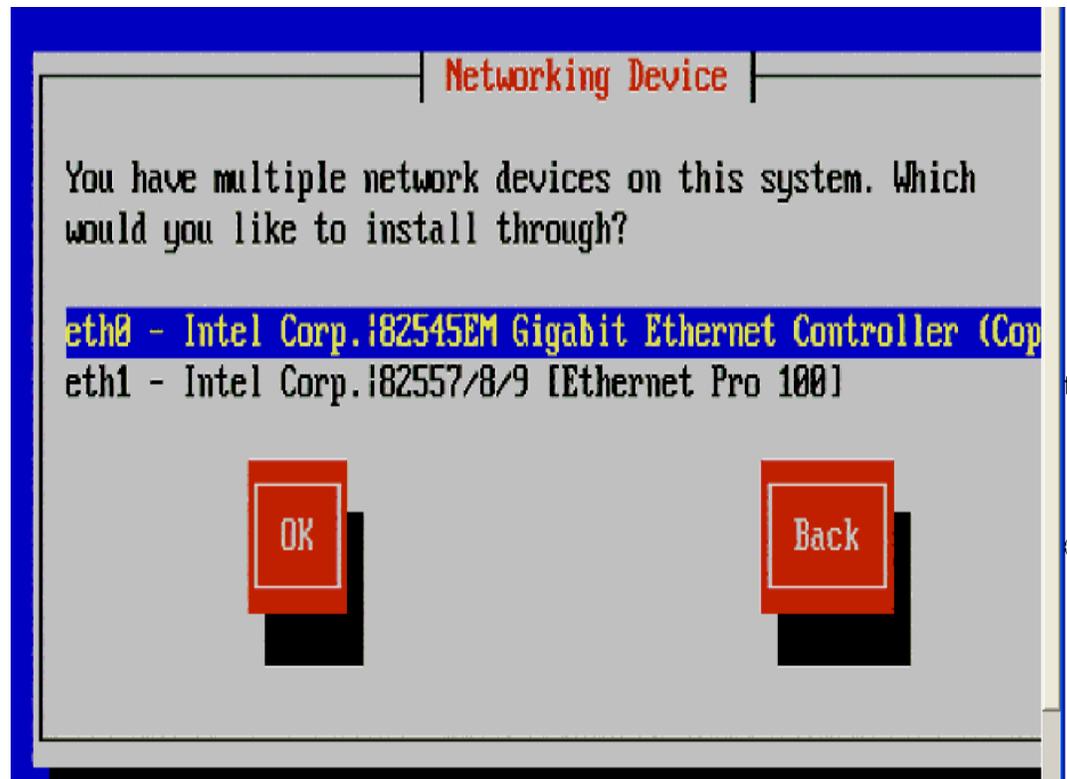


Figure 5-2: Installation Process

4. (RedHat only) At the **Do you have a driver disk?** prompt, enter Yes.

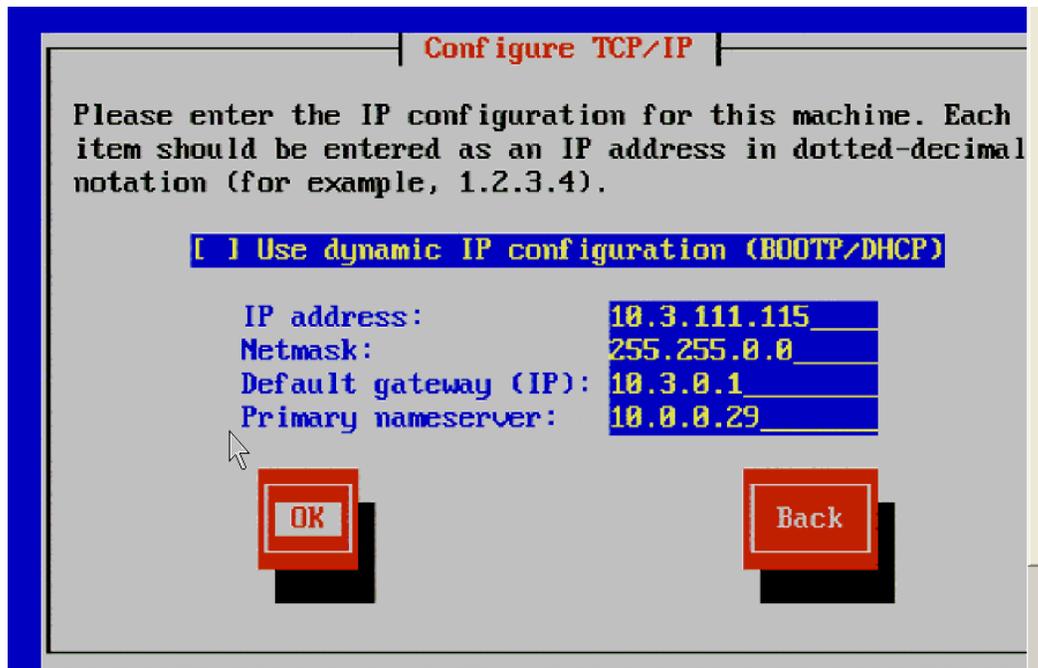


Figure 5-3: The Driver Disk Window

5. (RedHat only) On the **Driver Disk Source** screen, select your CD ROM, then select **OK**. The **Insert Driver Disk** screen appears.

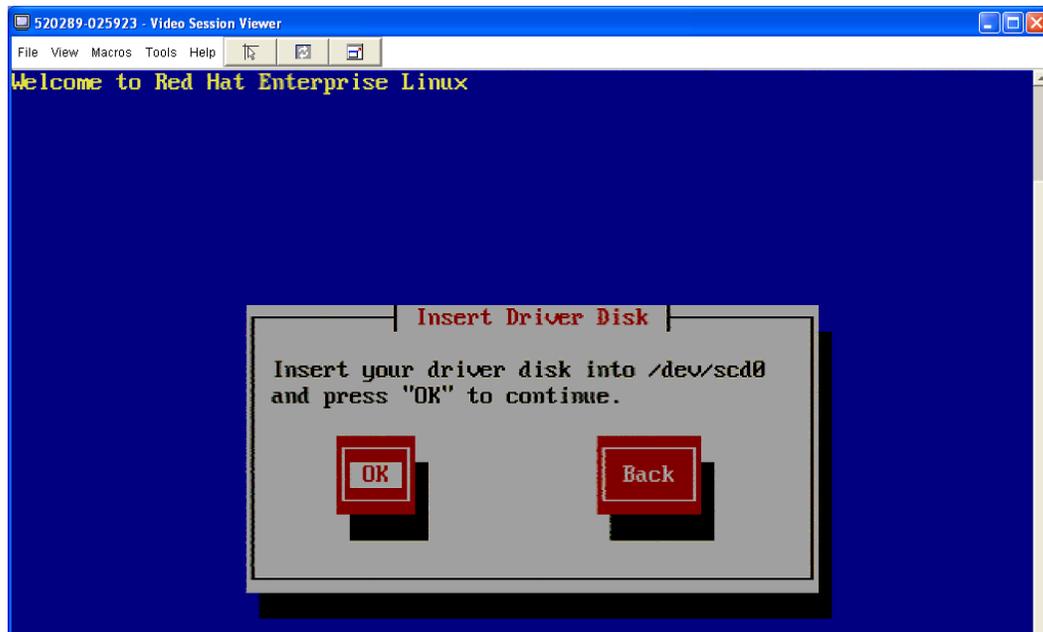


Figure 5-4: Insert the Driver Disk Window

- (RedHat only) Remove the Linux installation CD and place the Topspin driver CD in the drive, then select **OK**. The `ts_srp_host` driver loads.



Figure 5-5: Driver Disk Loading

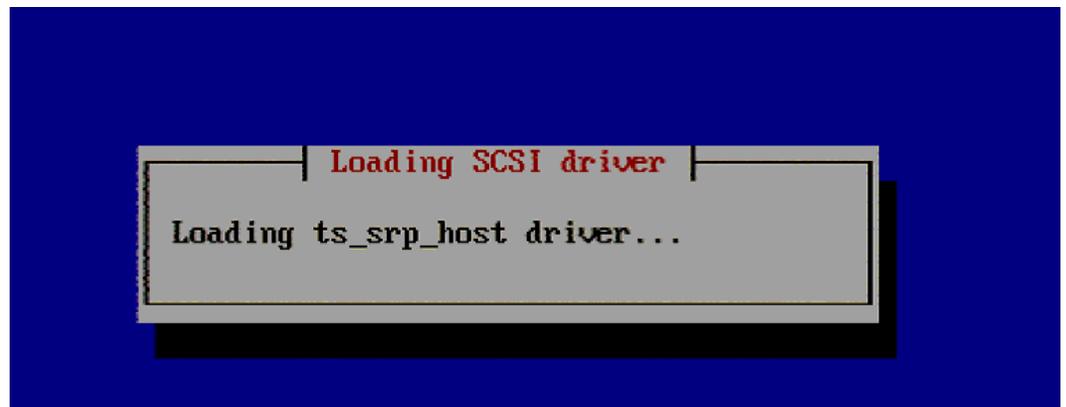


Figure 5-6: Loading SRP Host Driver

NOTE: Install the boot-loader to the SRP drive. If you have a local ATA drive, by default the host will be configured to install the boot-loader on the ATA drive, which you must avoid. Verify that the host installs the boot-loader on the SRP drive.

- (RedHat only) At the **Do you wish to load any more driver disks?** screen, select **no**.

The **Installation Method** window appears if you previously appended the **askmethod** option to the **linux dd** command (see [step 3](#)). If you did not use the **askmethod** option, this window will not appear. Continue to install from the CDs.

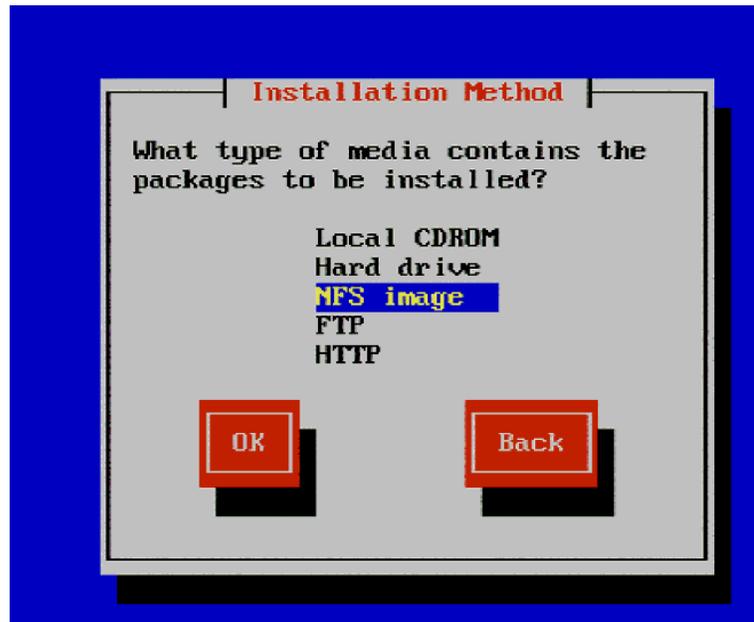


Figure 5-7: Select Installation Method

In this example we have chosen to complete the installation over NFS.



Figure 5-8: Select the Networking Device

8. Select the network interface over which you wish to install the image and click **OK**.

The **Configure TCP/IP** window appears.

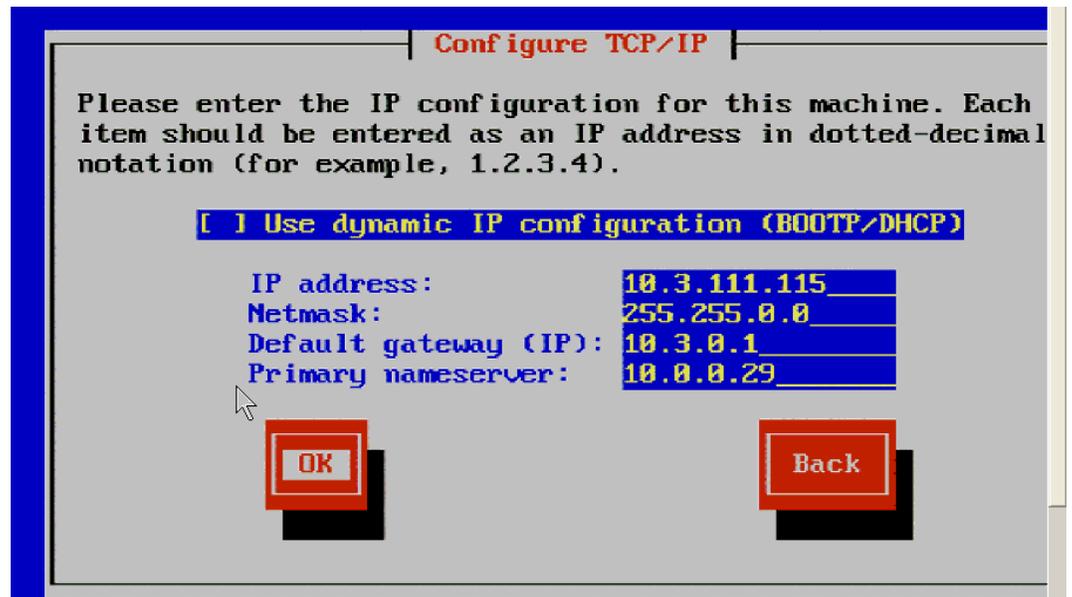


Figure 5-9: Enter the Network IP Configuration Information

9. Enter the appropriate IP Network information for the remote machine from which you wish to install the packages and click the **OK** button.

The **NFS Setup** window appears.

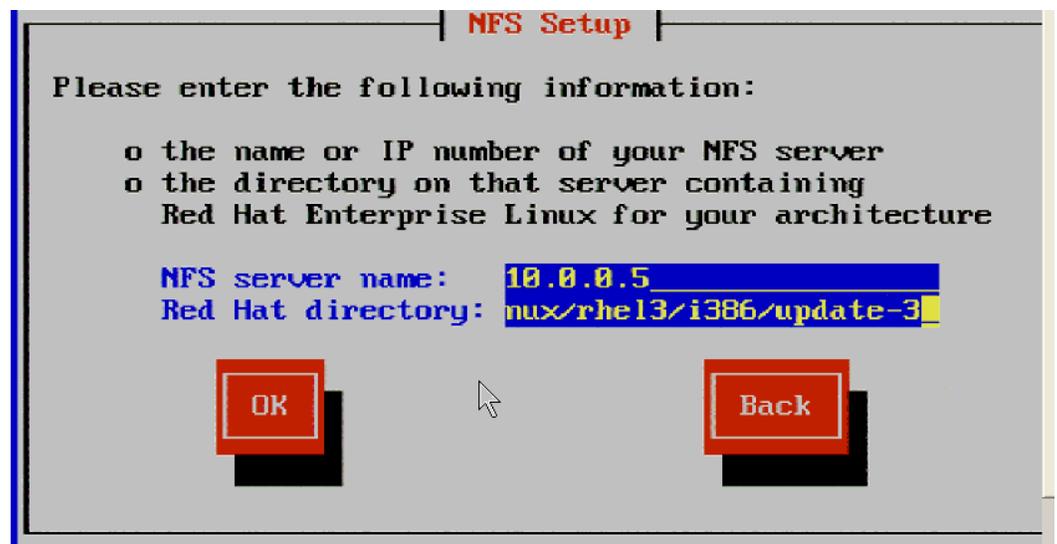


Figure 5-10: NFS Setup Information

10. Enter the **NFS server name** and the path to the **Red Hat directory** and click the **OK** button.

11. (RedHat only) Complete the installation as you normally would.

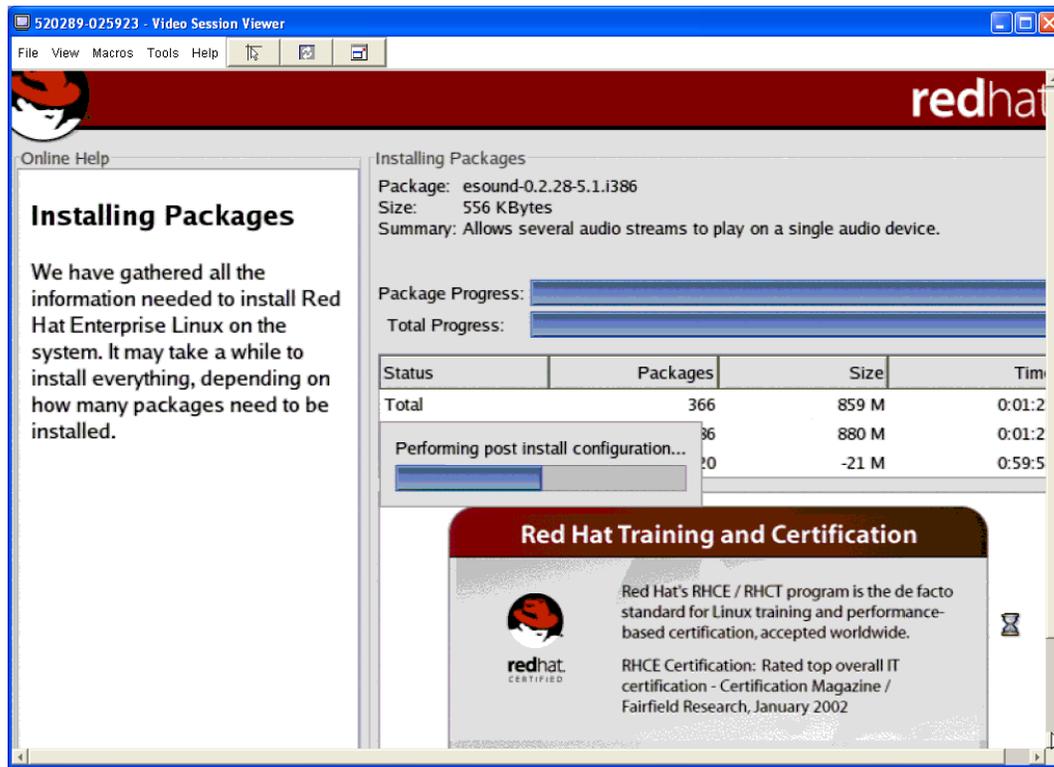


Figure 5-11: Installing Red Hat Packages

12. (RedHat only) After the installation is complete, remove the driver CD and reboot the host.
 13. Change the default boot order to *well known* so that SRP Boot is selected first.

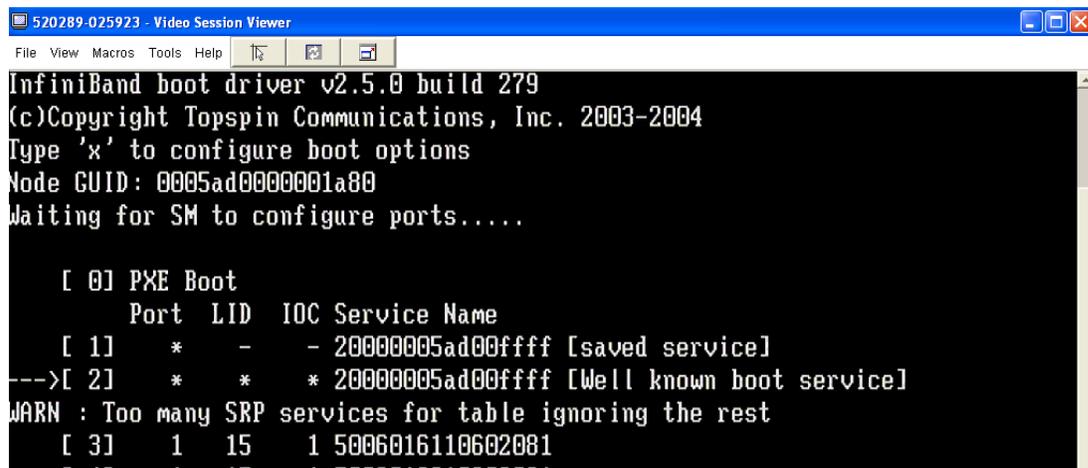


Figure 5-12: Select Well Known Boot Service

Use the `tvflash -o boot_type=well_known` command to change the default boot option without using the boot ROM itself.

Setting the value to well known requires the use of a boot ROM that is SRP capable.

Example

```
[root@ home]# /usr/local/topspin/sbin/tvflash -o boot_type=well_known
Writing [=====]
[root@ home]# /usr/local/topspin/sbin/tvflash -o
  auto_upgrade=yes
  boot_enable_port_1=yes
  boot_enable_port_2=yes
  boot_service_scan=yes
  boot_type=well_known
  boot_saved_port=0
  boot_saved_ioc_num=0
  boot_saved_dgid=0000:0000:0000:0000:0000:0000:0000:0000
  boot_saved_service_name=0000:0000:0000:0000
```



NOTE: If you do not install over NFS, you will be prompted to re-insert the Topspin CD at the end of the process.

Install and Boot Red Hat Enterprise Linux 2.1

Install with a local drive or (recommended) with no local drive.

No Local Drives

To configure the Boot over IB feature on the RHEL 2.1 and RHAS 2.1 platforms, perform the following steps:

1. Review the instructions from [“Configure a Fibre Channel Connection”](#) on page 72.
2. Restrict global SRP access as per [step 1](#).
3. Configure an ITL as per [step 2](#).
4. Connect a portable CD device to a diskless server. We recommend using a KVM device to view the installation process.
5. Place the first installation CD in your host and boot the server from the CD.
6. (RedHat only) At the **boot** prompt, enter the **linux dd** command to prompt the installer to run the Topspin driver CD.



NOTE: Advanced users can copy only RHAS 2.1 files, but you must maintain the directory structure.

7. Edit your kickstart file (configuration file). Add this line to your kickstart file on the NFS server:
driverdisk ip-address-of-host:/directory-containing-Topspin-BoIB-CD-contents --type nfs

Example

```
driverdisk 10.0.0.5:/data/home/rob/ks/dd --type nfs
```



NOTE: The kickstart file must reside on the NFS host that you enter during the install. Customize with the IP address of the host.

8. Update the kickstart file with the IP address, subnet mask, gateway, and host name of your SRP host.
9. Place disk 1 of RHAS 2.1 update 3 in your CD-ROM drive.
10. Boot your host and press the **x** key when the **Type ‘x’ to configure boot options** prompt appears during the boot process, then verify that your boot target appears in the list.
11. Press **x** to continue booting.
12. At the ISOLINUX **boot:** prompt, enter **linux ks=nfs:nfs-ip-address:/directory/kickstart-file**

Example

```
linux ks=nfs:10.0.0.5:/data/home/rob/as21.cfg
```

The host installs to the remote storage.

13. Reboot and enter the BIOS. Verify that “SRP boot” is first in the boot order. If it is not first, make it first in the boot order.

With Local Drives

To configure the Boot over IB feature on the RHEL 2.1 and RHAS 2.1 platforms, perform the following steps:

1. Review the instructions from “[Configure a Fibre Channel Connection](#)” on page 72.
2. Restrict global SRP access as per [step 1](#).
3. Configure an ITL as per [step 2](#).
4. (Select motherboards only) Download and compile a newer grub. For relevant motherboards, refer to release notes.
5. Copy the complete contents (entire hierarchy) of the Topspin Boot over IB CD to a directory on your NFS server.



NOTE: Advanced users can copy only RHAS 2.1 files, but you must maintain the directory structure.

6. Edit your kickstart file (configuration file). Add this line to your kickstart file on the NFS server:
driverdisk ip-address-of-host:/directory-containing-Topsin-BoIB-CD-contents --type nfs

Example

```
driverdisk 10.0.0.5:/data/home/rob/ks/dd --type nfs
```



NOTE: The kickstart file must reside on the NFS host that you enter during the install. Customize with the IP address of the host.

7. Update the kickstart file with the IP address, subnet mask, gateway, and host name of your SRP host.
8. Place disk 1 of RHAS 2.1 update 3 in your CD-ROM drive.
9. Boot your host and press the **x** key when the **Type ‘x’ to configure boot options** prompt appears during the boot process, then verify that your boot target appears in the list.
10. Press **x** to continue booting.
11. At the ISOLINUX **boot:** prompt, enter **linux ks=nfs:nfs-ip-address:/directory/kickstart-file**

Example

```
linux ks=nfs:10.0.0.5:/data/home/rob/as21.cfg
```

The host installs to the remote storage.

12. Reboot and enter the BIOS. Verify that “SRP boot” is first in the boot order. If it is not first, make it first in the boot order. Proceed with the SRP boot.
13. Log in to your host.
14. Edit `/etc/modules.conf` and change the aliases so that
 - the `ts_srp_host` is aliased to `scsi_hostadapter`
 - any additional SCSI controllers are aliased to `scsi_hostadapter1`, `2`, `3`, etc.
15. Enter the `/boot` directory, note the name of the `initrd` file, then rename or move the existing `initrd` file.
16. Enter the `mkinitrd` command with the original file name of the file that you moved in [step 15](#).

Example

```
mkinitrd initrd-2.4.9-e.34smp.img 2.4.9-e.34smp.
```



NOTE: If you change the order of drivers in `modules.conf`, you want to rebuild your `initrd` to ensure that your drivers load in the correct order.

17. Power down your host and re-insert your local drive, then boot your host.



NOTE: You may have to repeat [step 12](#). Some operating systems place new devices at the top of the boot order automatically.



NOTE: To ensure a successful boot, install with a separate `/boot` partition.



NOTE: If you do not install over NFS, you will be prompted to re-insert the Topspin CD at the end of the process.

Install and Boot SUSE Linux Enterprise Server 8 Service Pack 3

To install SLES 8, Service Pack 3, you *must insert CD-ROM disks in the correct order*. Pay strict attention to the instructions that follow to successfully boot. To install SLES 8 Service Pack 3 to boot over SRP, perform the following steps:

1. Collect the SUSE installation CDs, Topspin drivers CD, and Service Pack CDs.



NOTE: NFS installs are not available for SUSE.

2. Place the Service Pack CD 1 in your CD-ROM drive and boot your host.

3. Select **OK** and prepare to press the **Alt** key when the boot prompt screen appears (Figure 5-13).

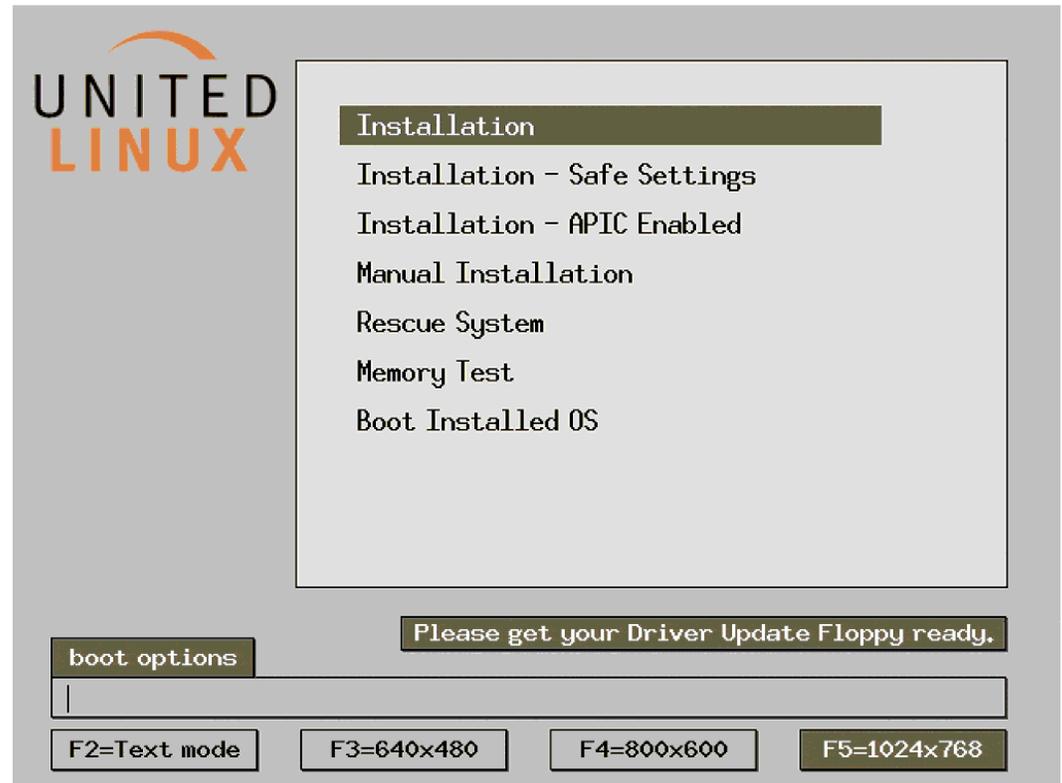


Figure 5-13: Boot Prompt

4. Press the **Alt** key when the boot prompt screen appears (Figure 5-13).
5. Highlight **Installation** in the menu.

6. Press the **Enter** key to load linux. A loading progress bar appears (Figure 5-14).

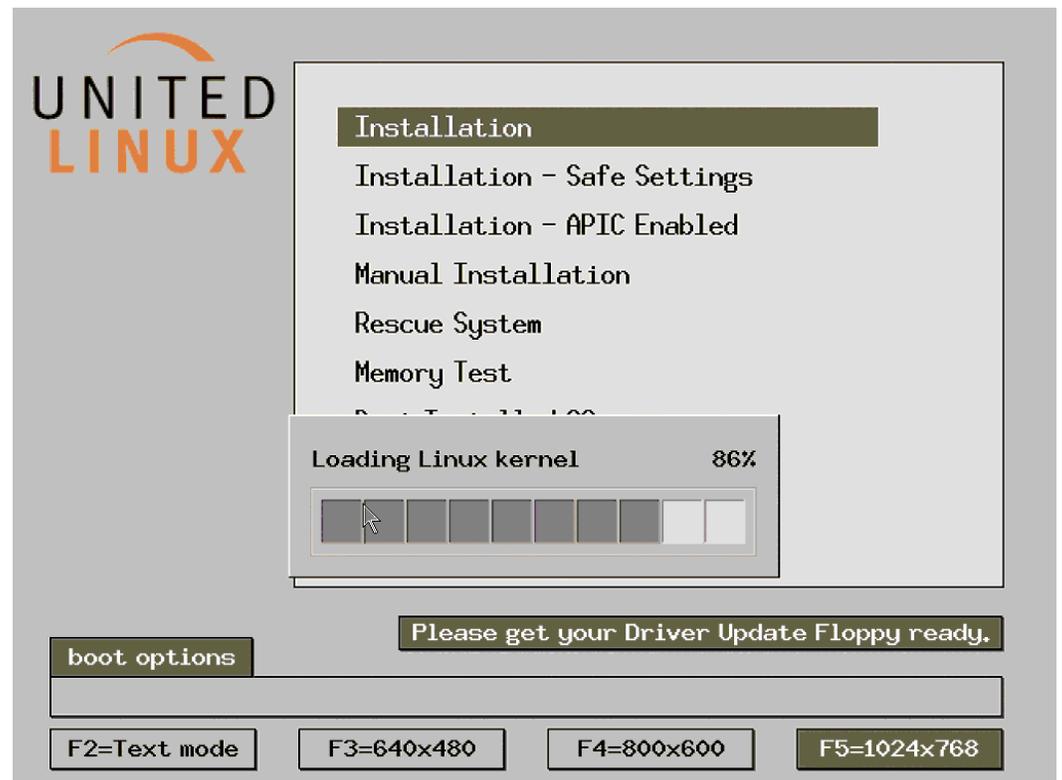
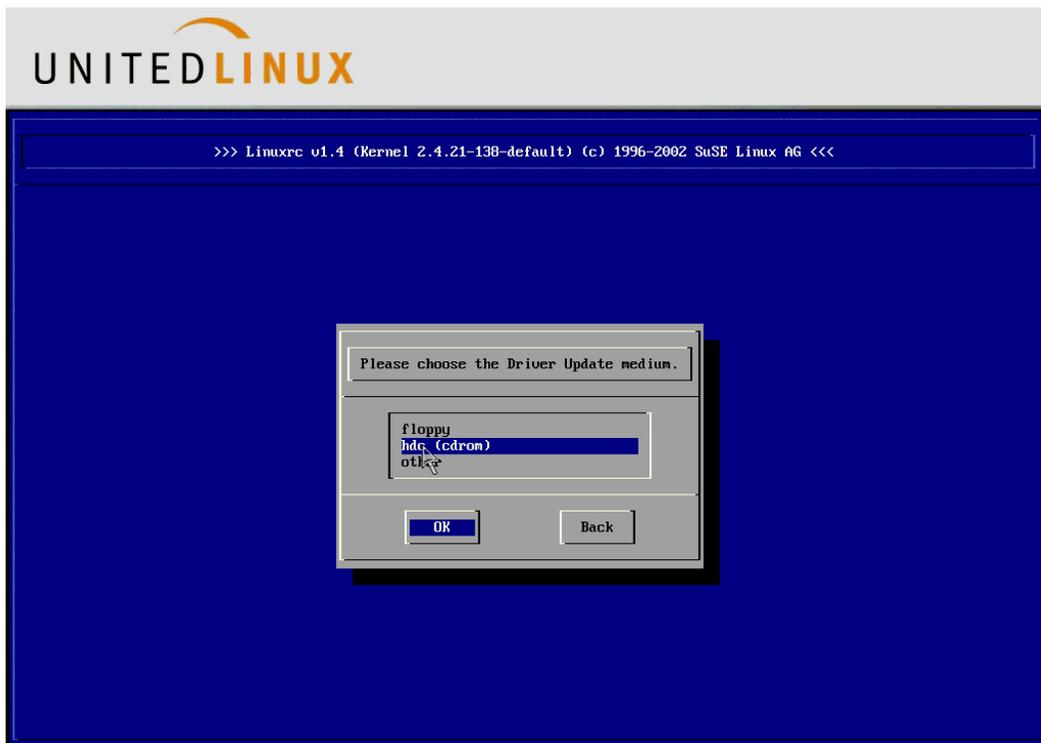


Figure 5-14: Progress Bar

7. At the “Please insert” prompt, remove the Service Pack CD 1 disk and place the Topspin Boot over InfiniBand CD in your host. The **Please choose the Driver Update medium** prompt appears.

8. Select **cdrom** at the **Please choose the Driver Update medium** prompt.

Example

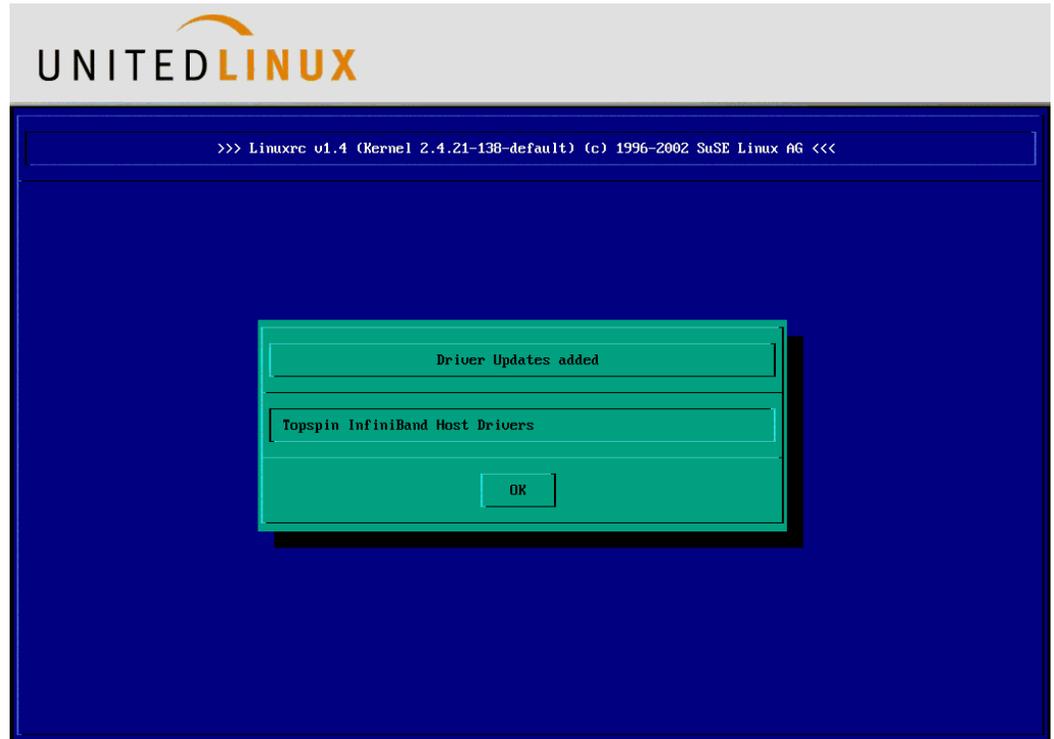


9. Press the **Enter** key. Topspin drivers begin to load.

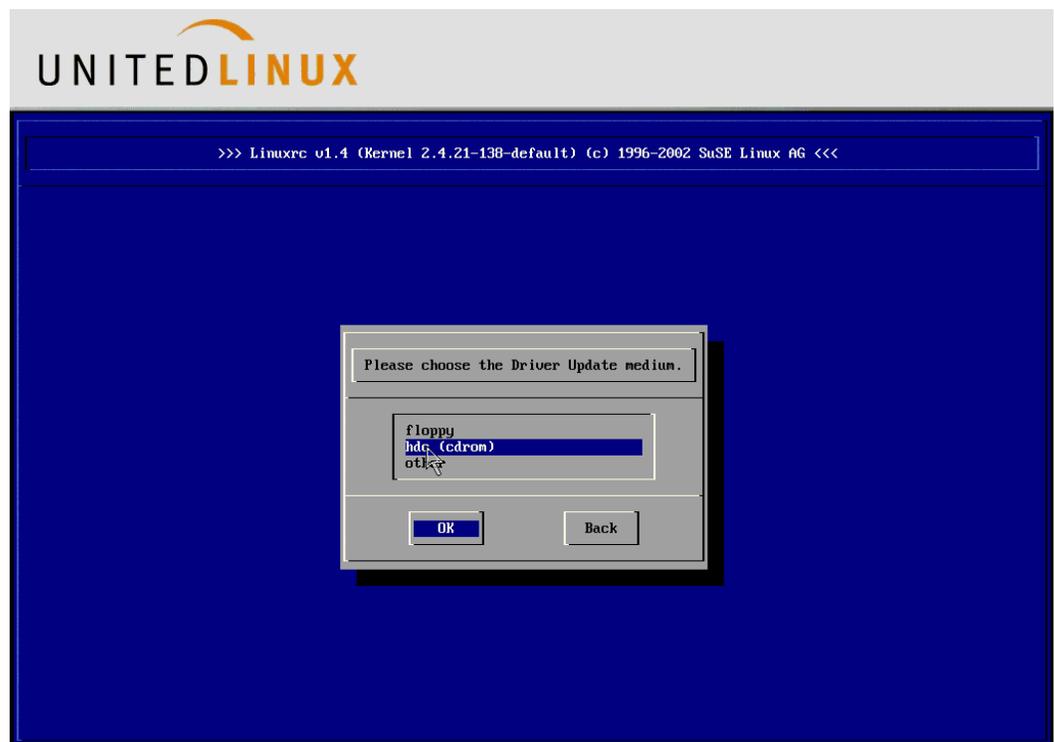


NOTE: To view installation details, press and hold **Ctrl-Alt-F3** to view output and **Ctrl-Alt-F4** to view the console. (See Linux distribution documentation for details.) Press and hold **Ctrl-Alt-F1** to return to primary display.

- Click **OK** after drivers load.

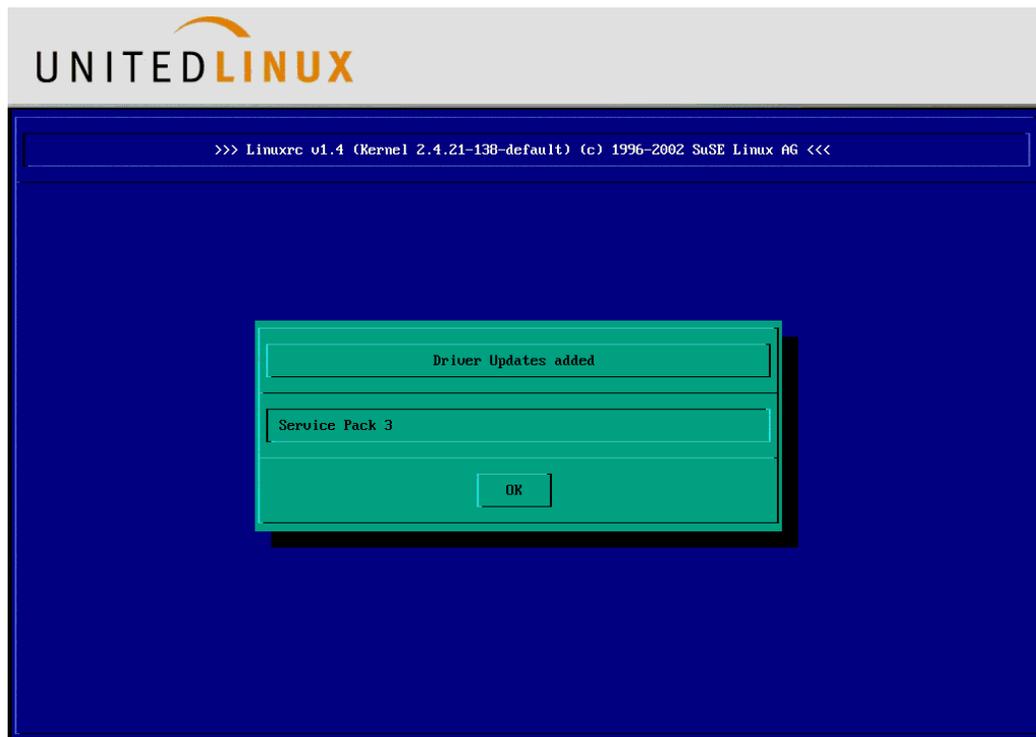
Example

- Return Service Pack CD 1 into drive.
- Select **OK** in the **Please Choose the Driver Update medium** window. The driver updates added window appears.

Example

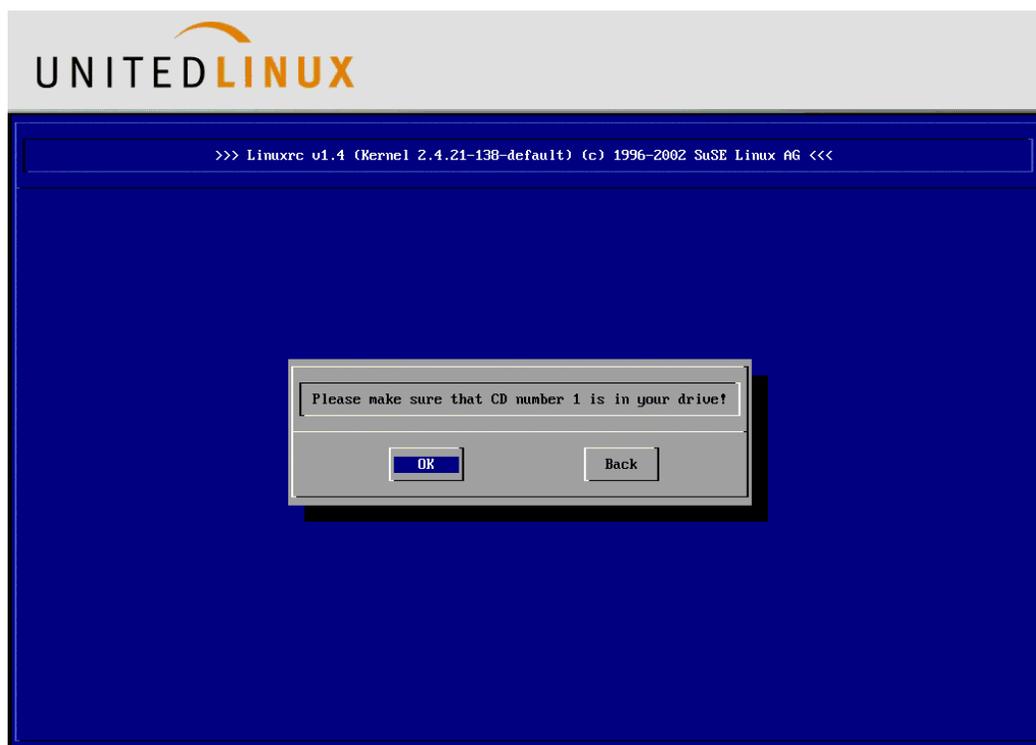
13. Select **OK**.

Example



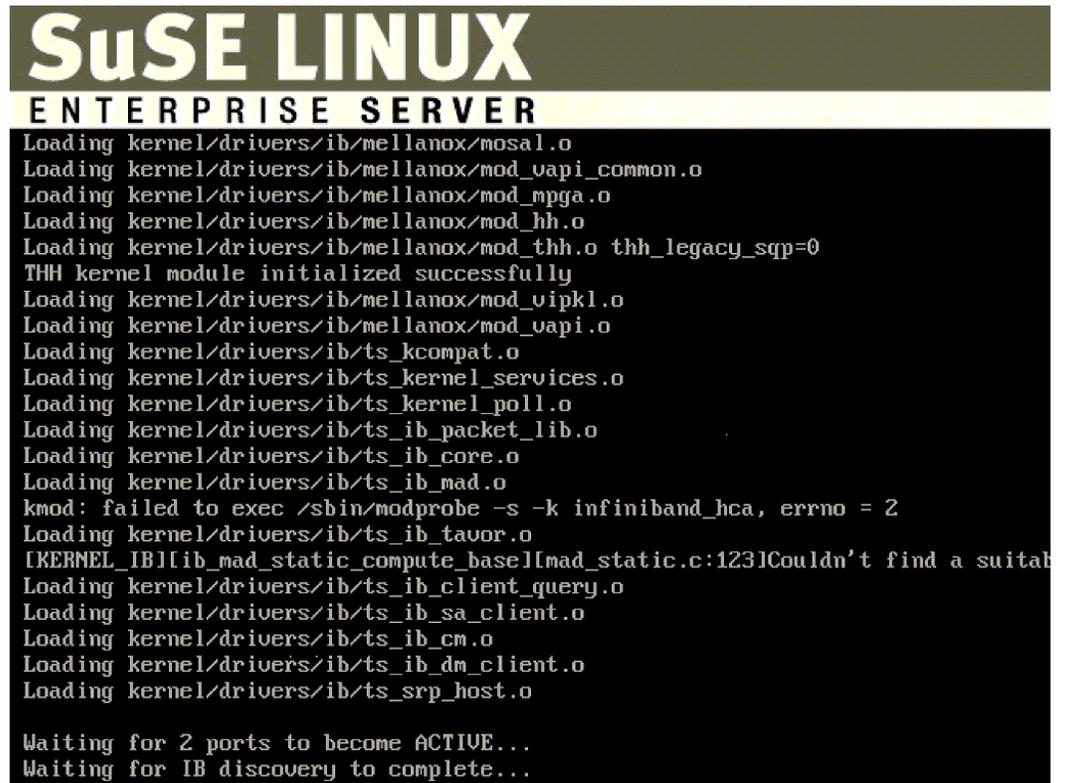
14. Select **Back** and press the **Enter** key.

Example



15. Put the *original* (not Service pack) SLES CD 1 in your CD-ROM drive.

16. Select **OK**. As your installation completes, you will see the Topspin drivers load (Figure 5-15) if you press and hold **Ctrl-Alt-F3** to view the installation.



```

SuSE LINUX
ENTERPRISE SERVER
Loading kernel/drivers/ib/mellanox/mosal.o
Loading kernel/drivers/ib/mellanox/mod_vapi_common.o
Loading kernel/drivers/ib/mellanox/mod_mpga.o
Loading kernel/drivers/ib/mellanox/mod_hh.o
Loading kernel/drivers/ib/mellanox/mod_thh.o thh_legacy_sqp=0
THH kernel module initialized successfully
Loading kernel/drivers/ib/mellanox/mod_vipkl.o
Loading kernel/drivers/ib/mellanox/mod_vapi.o
Loading kernel/drivers/ib/ts_kcompat.o
Loading kernel/drivers/ib/ts_kernel_services.o
Loading kernel/drivers/ib/ts_kernel_poll.o
Loading kernel/drivers/ib/ts_ib_packet_lib.o
Loading kernel/drivers/ib/ts_ib_core.o
Loading kernel/drivers/ib/ts_ib_mad.o
kmod: failed to exec /sbin/modprobe -s -k infiniband_hca, errno = 2
Loading kernel/drivers/ib/ts_ib_tavor.o
[KERNEL_IB]ib_mad_static_compute_base[ib_mad_static.c:123]Couldn't find a suitable
Loading kernel/drivers/ib/ts_ib_client_query.o
Loading kernel/drivers/ib/ts_ib_sa_client.o
Loading kernel/drivers/ib/ts_ib_cm.o
Loading kernel/drivers/ib/ts_ib_dm_client.o
Loading kernel/drivers/ib/ts_srp_host.o

Waiting for 2 ports to become ACTIVE...
Waiting for IB discovery to complete...

```

Figure 5-15: Topspin Drivers



NOTE: When you partition in the standard setup, we recommend that you create a separate /boot partition.

17. Place Service Pack 3 CD 1 in your host when your installation completes to install additional updates.



NOTE: You must perform this process twice to confirm that you have the correct modutils.

To install additional updates, perform the following steps:

- a. Boot your host.
 - b. Place the SLES 8 Service Pack 3 CD 1 in your CD-ROM drive.
 - c. Mount your CD-ROM drive.
 - d. Launch YaST2.
 - e. Double-click **Patch CD Update**. The update display launches.
 - f. Click **Expert**.
 - g. Click the **CD** radio button.
 - h. Click the radio button of the CD-ROM drive that you mounted in [step c](#).
 - i. Click **OK**.
 - j. Click **Manual Update**.
 - k. Click **Next**.
 - l. Check the **Recommended updates for YOU** checkboxes.
 - m. Click **Accept**.
 - n. Click **OK** as the install proceeds.
 - o. Click **Next** when the install completes.
18. Repeat [step 17](#).



NOTE: As of this release, you must perform this operation twice to ensure that the updates install successfully.

19. Confirm that your host now runs the appropriate modutils.
- a. Open a command-line window.
 - b. Enter the **rpm -q modutils** command and verify that 2.4.25-53 appears in the output.

Install and Boot SUSE Linux Enterprise Server 9 Service Pack 1



NOTE: This distribution can be installed directly from the Topspin BoIB CD. Older supported SLES9 (currently FCS) are available, but as a smaller ISO image that is included on our BoIB CD. You have to burn the smaller ISO image to a CD yourself.

To install SLES 9, Service Pack 1, on a remote SRP disk, perform the following steps:

1. Insert the Service Pack CD 1 in your host.
2. Boot your host. Prepare to press the F6 key when the screen in [Figure 5-16](#) appears.

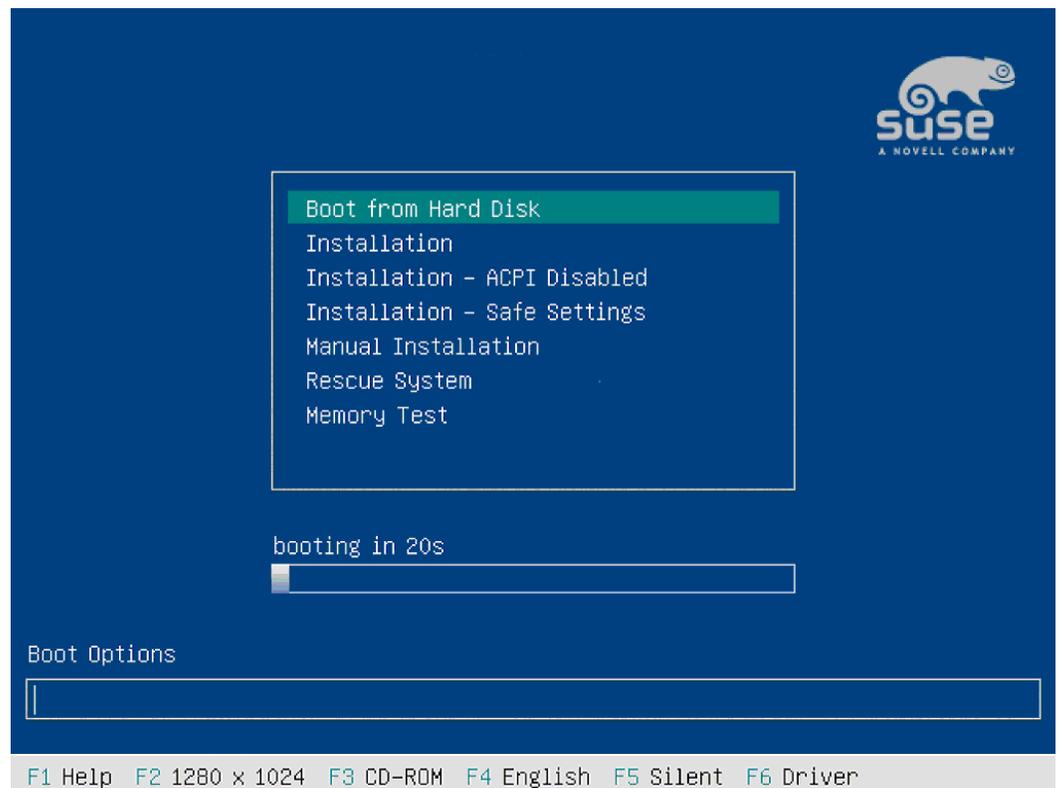
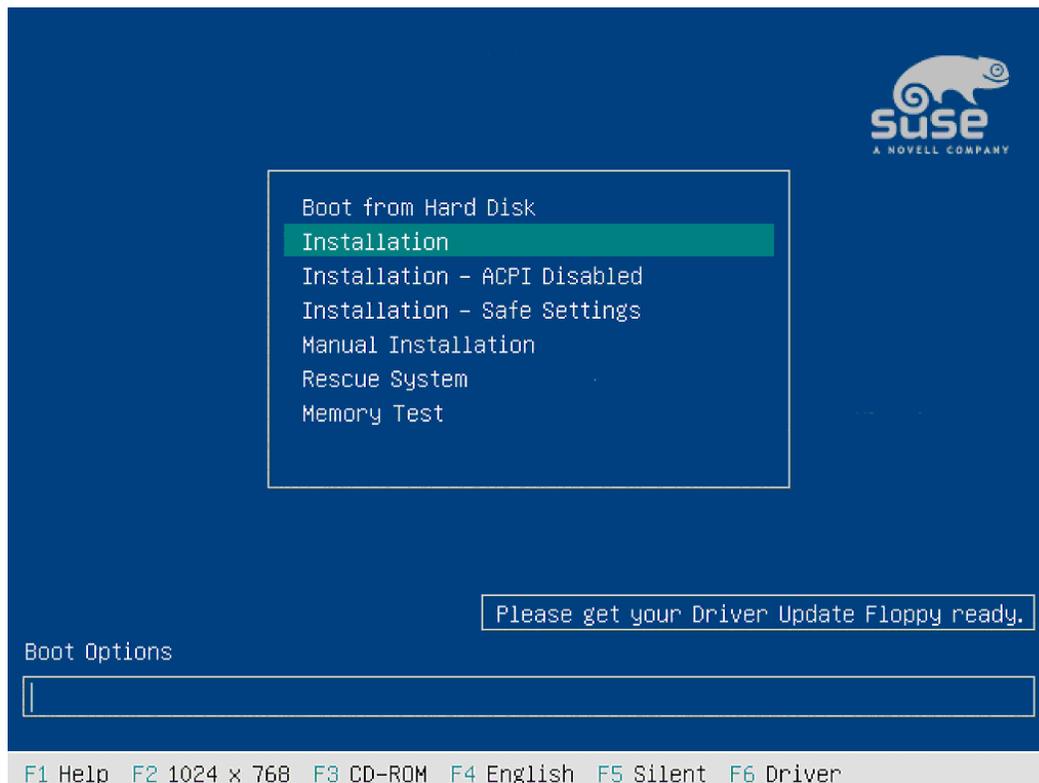


Figure 5-16: Initial GUI Screen

3. Press the **F6** key.

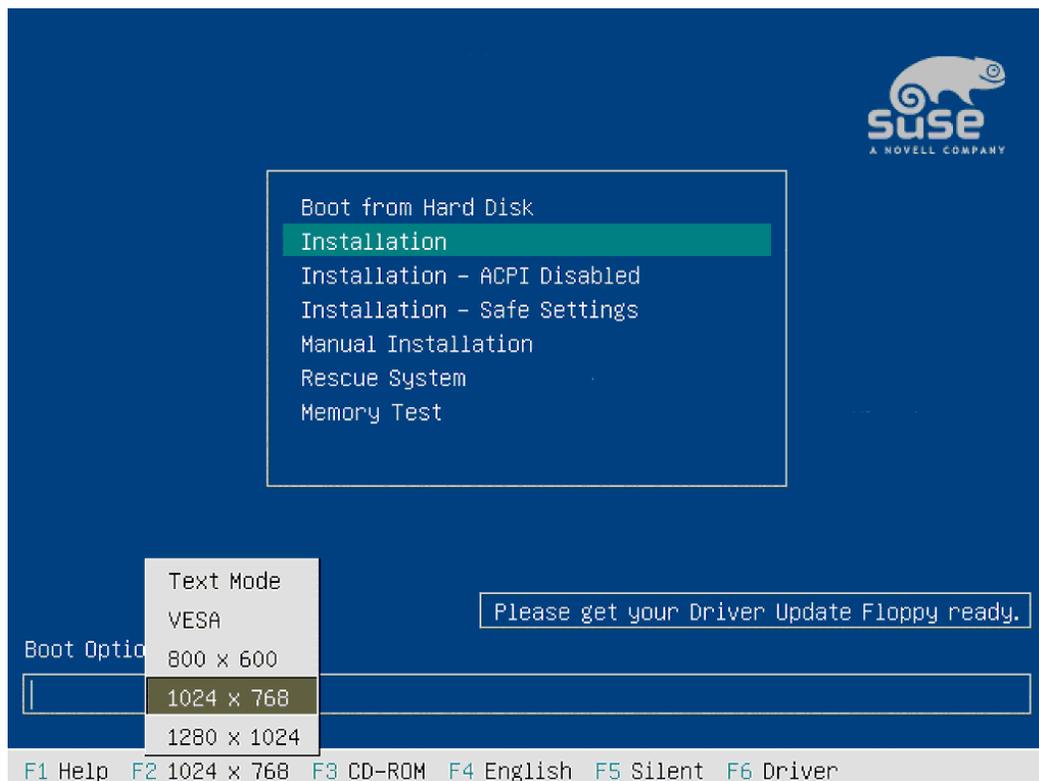
- Highlight **Installation** in the GUI.

Example



- Press the **F2** key and highlight 1024 x 768 to configure the screen resolution.

Example



6. Press the **Enter** key to continue the installation. The screen in [Figure 5-17](#) appears.

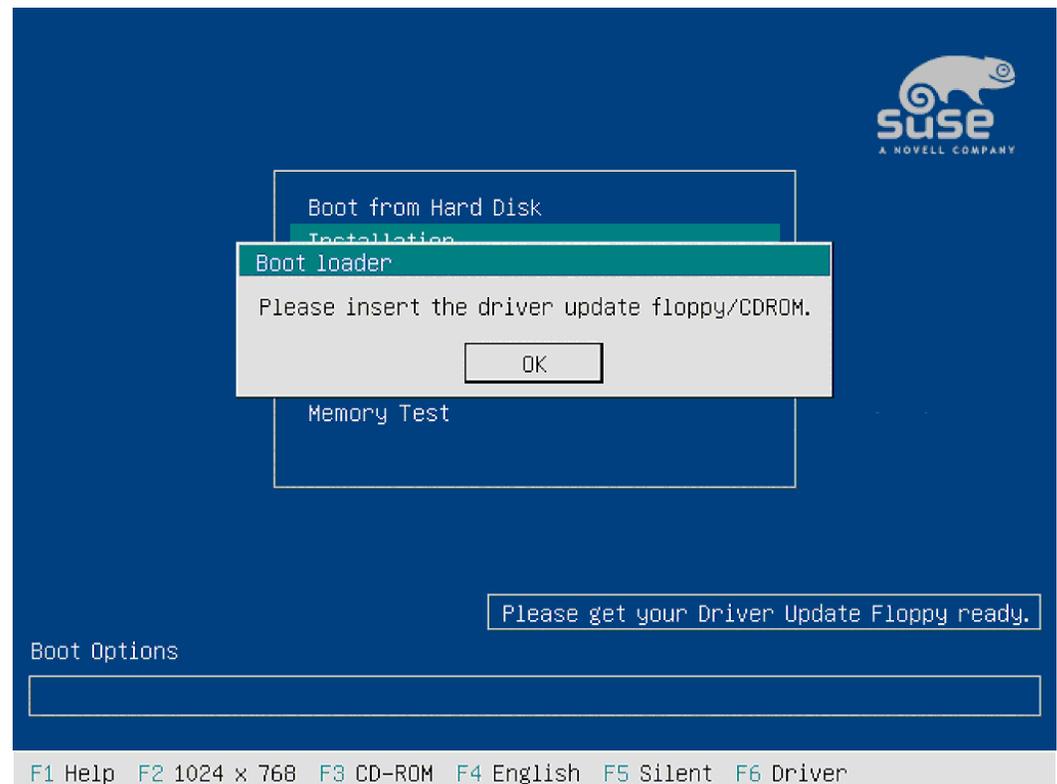
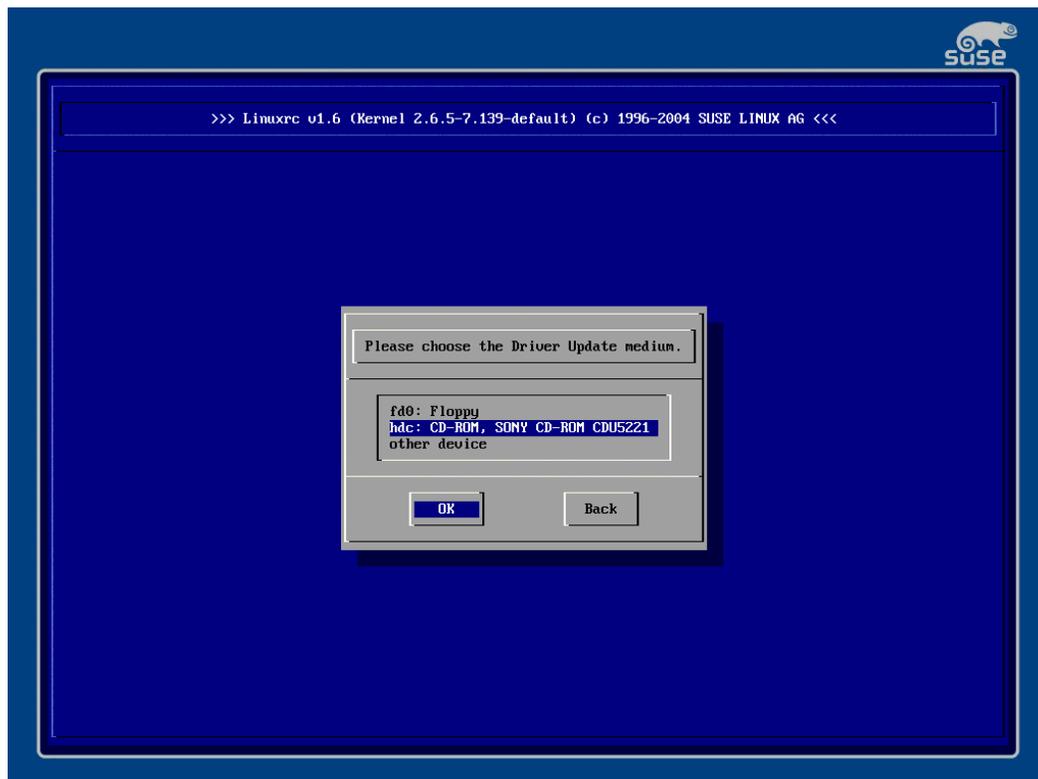


Figure 5-17: Disk-Change Prompt.

7. Place your Topspin drivers CD in your host.
8. Press the **Enter** key to continue.

9. Highlight **CD-ROM**.

Example



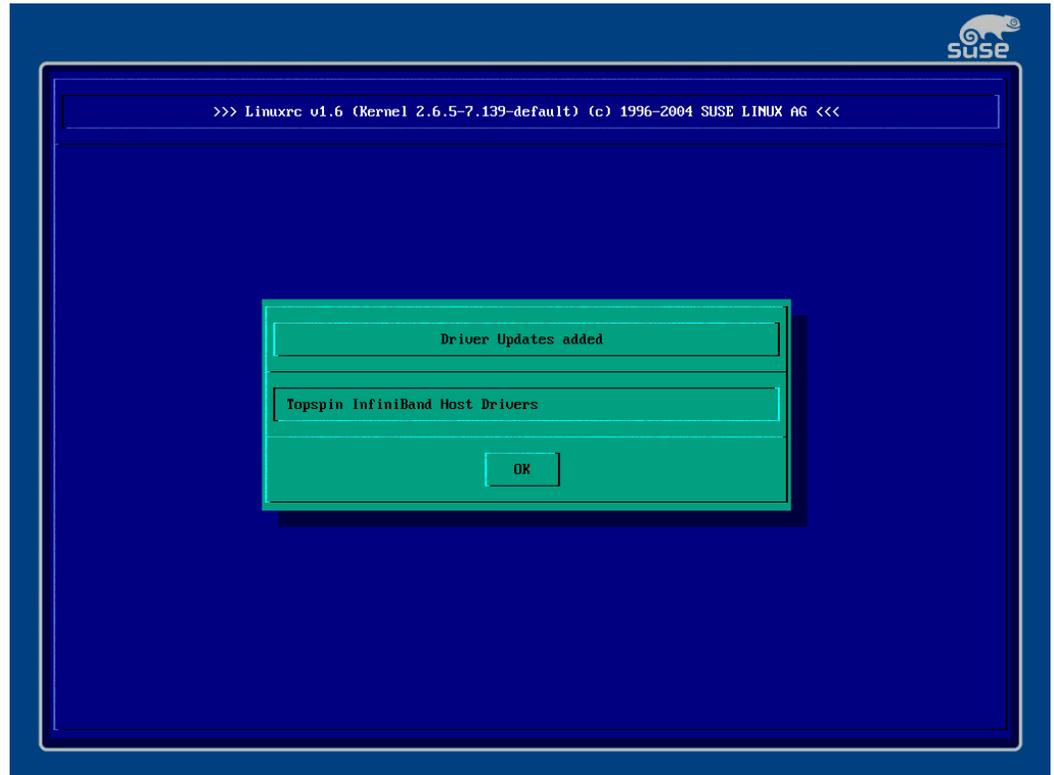
10. Highlight **OK** and press the **Enter** key.



NOTE: To view installation details, press and hold **Ctrl-Alt-F3** to view output and **Ctrl-Alt-F4** to view the console. (See Linux distribution documentation for details.) Press and hold **Ctrl-Alt-F1** to return to primary display.

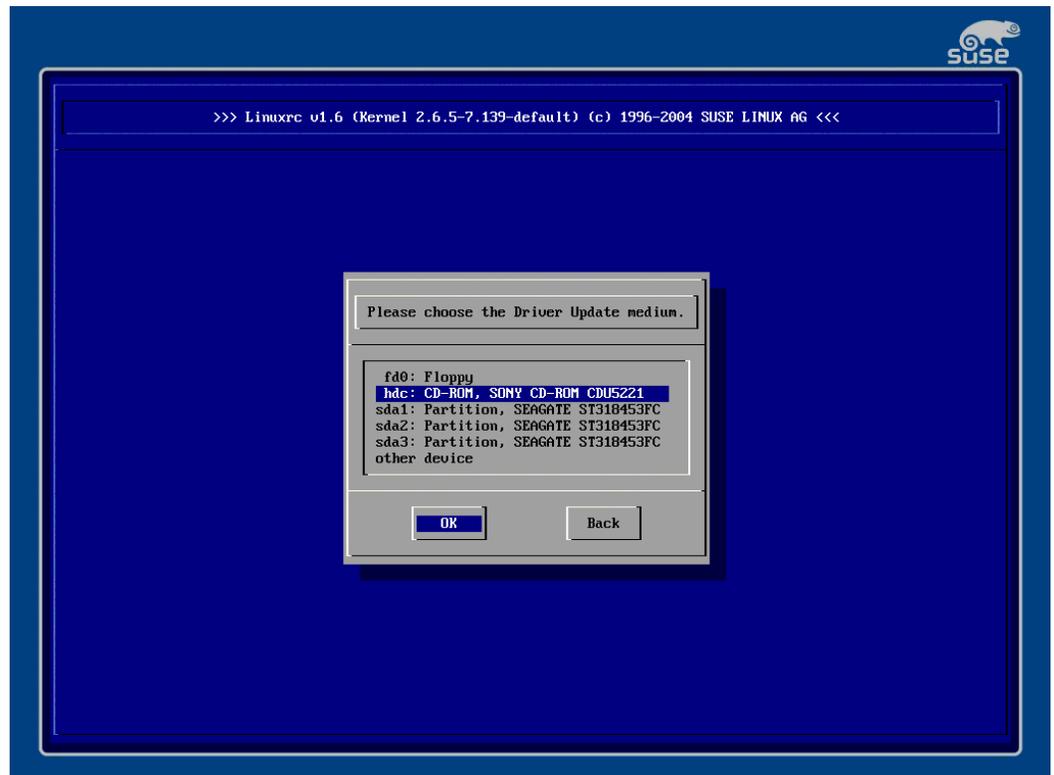
11. Select **OK** and press the **Enter** key.

Example



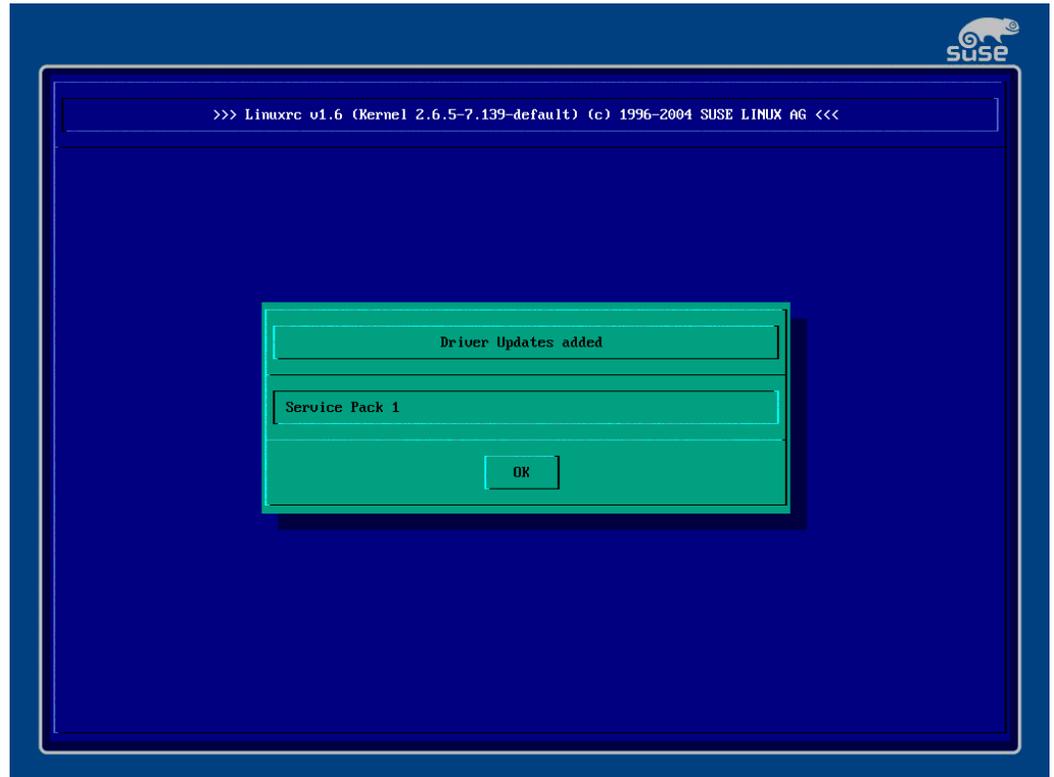
12. Place the service pack CD 1 back in your host.
13. Select **OK** and press the **Enter** key.

Example



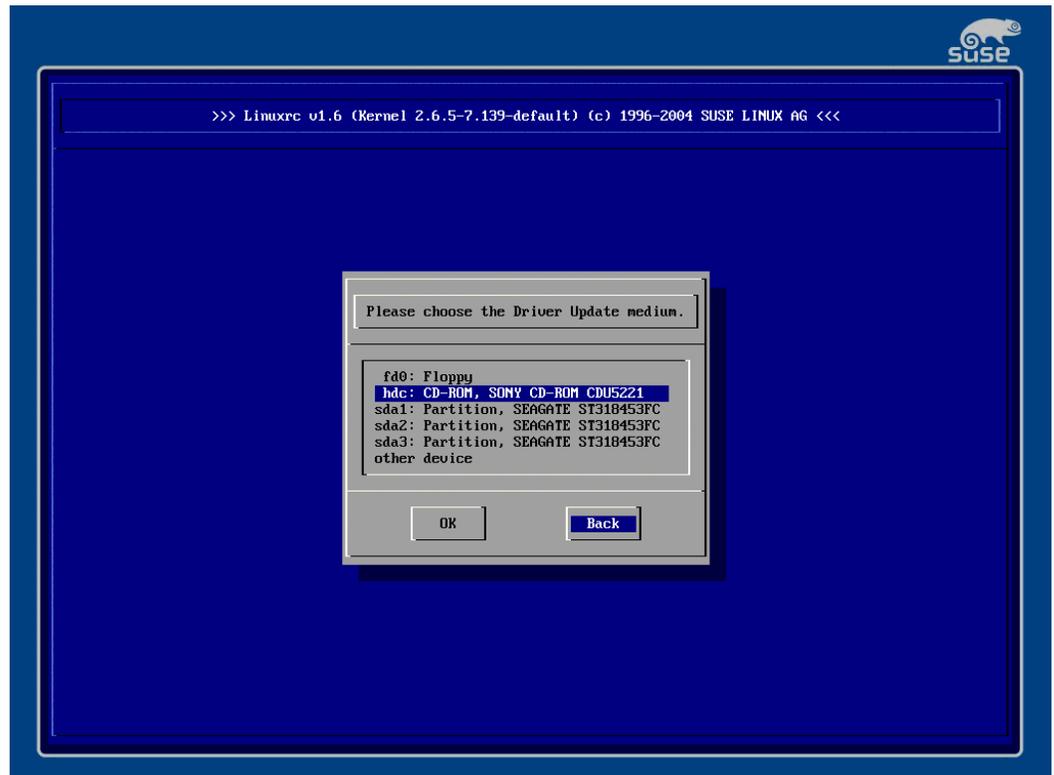
14. Select OK and press the Enter key.

Example



15. Select **Back** and press the **Enter** key.

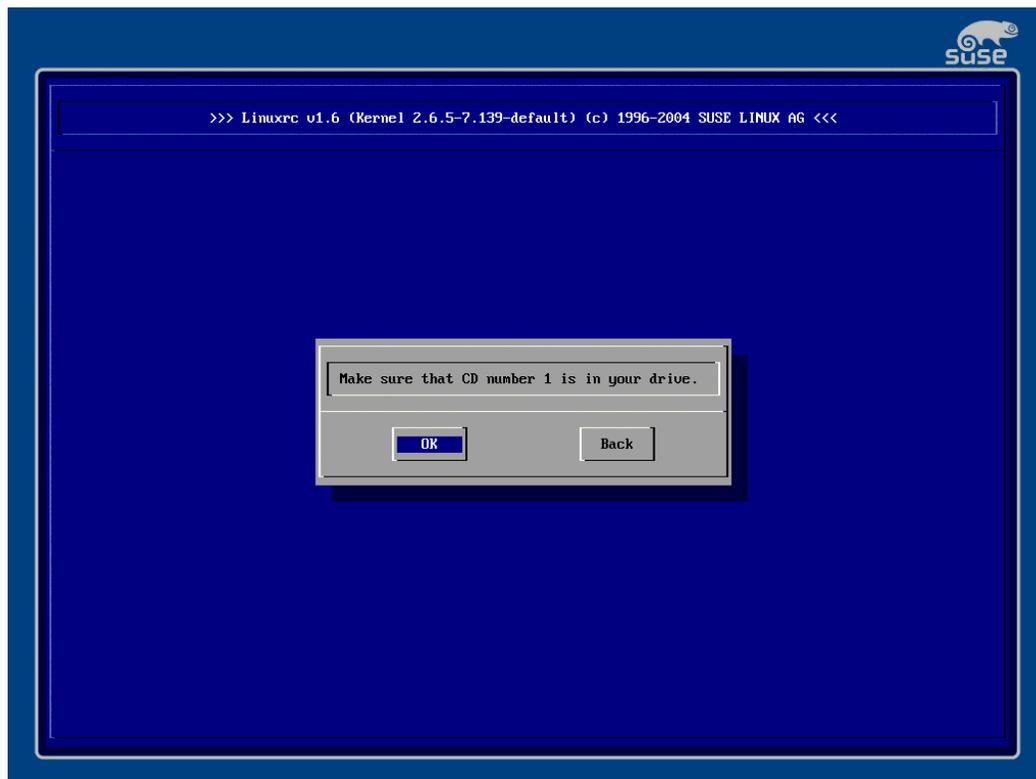
Example



16. Place the SLES 9 CD 1 in your host.

17. Select **OK** and press the **Enter** key.

Example



NOTE: To ensure a successful boot, install with a separate /boot partition during the standard Linux install steps.

Boot an Existing Linux Image

After you install an image onto Fibre Channel storage with a host, you can continue to boot that host from the storage, or you can replace that host with a new, identical host.

Boot the Installation Host

To boot the remote image with your installation host, simply reload the host. If the boot fails, verify that you

- Reconfigured the boot order of the host to put SRP boot first.
- Reconfigured the boot service to use Well Known Service Name first and configured the WKNS on your Server Switch (see [step n](#) of “Use the CLI and a Host with a Remote Boot HCA” on page 73).

Boot a New, Diskless Host with a New HCA

The instructions in this section assume the following:

- Your host includes an HCA that runs Boot over IB firmware.
- You have configured the Fibre Channel gateway on your Server Switch to deny port and LUN access to all new initiators.

Use the CLI

To boot an existing image with a new host and HCA, perform the following steps:

1. Remove or reconfigure any host that runs the existing image. Only one host on your network may boot a given image.
2. Physically connect your new host to your Server Switch.
3. Configure an ITL.



NOTE: To install an image with VFrame, do not configure ITLs manually.

- a. Open your host and record the GUID from the label on your HCA. Alternatively, you can boot your host and press the **x** key when the **Type 'x' to configure boot options** prompt appears during the boot process. The GUID of the host appears in the host display. Record the GUID and leave the host at this screen. You will refresh the screen later in this process.
- b. Log into your Server Switch.
- c. Enter the **enable** command to enter Privileged Exec mode.

Example

```
Topspin-360> enable
Topspin-360#
```

- d. Enter the **configure terminal** command to enter Global Configuration mode.

Example

```
Topspin-360# configure terminal
Topspin-360 (config) #
```

- e. Enter the **fc srp initiator** command with
 - the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
 - the **auto-bind** keyword
 to add your host to the configuration file and create WWPNs, on the Fibre Channel gateway, that point to the host.

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 auto-bind
```

- f. (Optional) Enter the **fc srp initiator** command with
 - the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
 - the **description** keyword
 - a text description, enclosed in quotation marks
 to assign an easily-recognizable identifier to the initiator.

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 description "test-initiator"
```

- g. Enter the **exit** command to return to Privileged Exec mode.

Example

```
Topspin-360 (config) # exit
Topspin-360 #
```

- h. Enter the **show fc srp initiator** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
- to view the WWPNs of the FC gateway ports that point to the host.

Example

```
Topspin-360 # show fc srp initiator
```

```
=====
                        SRP Initiators
=====
      guid: 00:05:ad:00:00:00:17:3c
      extension: 00:00:00:00:00:00:00:00
      description: test-initiator
      wwnn: 20:02:00:05:ad:00:00:00
      credit: 0
      active-ports: none
      pkeys:
      bootup-target: 00:00:00:00:00:00:00:00
      bootup-lu: 00:00:00:00:00:00:00:00
      action: auto-bind
      result: success
      wwpns: port      wwpn                      fc-addr
              2/1      20:02:00:05:ad:20:00:00    00:00:00
              2/2      20:02:00:05:ad:24:00:00    00:00:00
              3/1      20:02:00:05:ad:30:00:00    00:00:00
              3/2      20:02:00:05:ad:34:00:00    00:00:00
              4/1      20:02:00:05:ad:40:00:00    00:00:00
              4/2      20:02:00:05:ad:44:00:00    00:00:00
              5/1      20:02:00:05:ad:50:00:00    00:00:00
              5/2      20:02:00:05:ad:54:00:00    00:00:00
              6/1      20:02:00:05:ad:60:00:00    00:00:1d
              6/2      20:02:00:05:ad:64:00:00    00:00:1d
              7/1      20:02:00:05:ad:70:00:00    00:00:00
              7/2      20:02:00:05:ad:74:00:00    00:00:00
              8/1      20:02:00:05:ad:80:00:00    00:00:00
              8/2      20:02:00:05:ad:84:00:00    00:00:00
              9/1      20:02:00:05:ad:90:00:00    00:00:00
              9/2      20:02:00:05:ad:94:00:00    00:00:00
             10/1      20:02:00:05:ad:a0:00:00    00:00:00
             10/2      20:02:00:05:ad:a4:00:00    00:00:00
             11/1      20:02:00:05:ad:b0:00:00    00:00:00
             11/2      20:02:00:05:ad:b4:00:00    00:00:00
             12/1      20:02:00:05:ad:c0:00:00    00:00:00
             12/2      20:02:00:05:ad:c4:00:00    00:00:00
             13/1      20:02:00:05:ad:d0:00:00    00:00:00
             13/2      20:02:00:05:ad:d4:00:00    00:00:00
```

- i. Record the WWPN of the FC gateway port that connects the host to the storage.
- j. Enter the **show fc srp initiator-wwpn-view** command with

- the WWPN of the port (from [step i](#))
- the **target** keyword

to verify that the host can see the target that contains the bootable image. If the WWPN of the boot target does not appear in the command output, give the gateway some time to configure the ITs and retry the command.

Example

```
Topspin-360# show fc srp initiator-wwpn-view 20:02:00:05:ad:60:00:00 target
```

```
=====
SRP Targets Accessible to Initiator Via Port WWN 20:02:00:05:ad:60:00:00
=====
          wwpn: 21:00:00:04:cf:f6:c2:ab
          wwnn: 20:00:00:04:cf:f6:c2:ab
description: SRP.T10:21000004CFF6C2AB
      ioc-guid: 00:05:ad:00:00:00:15:1a
service-name: SRP.T10:21000004CFF6C2AB
protocol-ids: 04:00:00:00:00:00:00:00:00:00
      fc-address: 00:00:23
          mtu: 0
connection-type: nl-port
physical-access: 6/1
```

- k. Enter the **configure terminal** command to return to Global Configuration mode.

Example

```
Topspin-360# configure terminal
Topspin-360 (config) #
```

- l. Enter the **no fc srp it** command with
 - the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
 - the WWPN of the target that stores the bootable image
 - the **gateway-portmask-policy** keyword
 - the **restricted** keyword
 - the FC gateway port that connects the host to the storage (in slot#/port# format) to grant the host access to the target from which your host will boot (through the port).

Example

```
Topspin-360 (config) # no fc srp it 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 21:00:00:04:cf:f6:c2:ab gateway-portmask-policy
restricted 6/1
```

- m. Enter the **fc srp initiator** command with
 - the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00:00**)
 - the **discover-itl** keyword

to configure ITLs to which the host belongs. The ITLs inherit the portmask policy of the IT (from [step l](#)).

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 discover-itl
```

- n. Enter the **no fc srp itl** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00**)
 - the WWPN of the target that stores the bootable image
 - the LUN ID of the LUN that stores the bootable image
 - the **lun-policy** keyword
 - the **restricted** keyword
- to grant the host access to the LUN from which your host will boot.

Example

```
Topspin-360 (config) # no fc srp itl 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 21:00:00:04:cf:f6:c2:ab 00:00:00:00:00:00:00
lun-policy restricted
```

- o. Enter the **fc srp initiator** command with
- the GUID of the host (from [step a](#))
 - the GUID extension of the host (always **00:00:00:00:00:00:00**)
 - the **bootup** keyword
 - the **target** keyword
 - the WWPN of the target that stores the boot image.
 - the **lu** keyword
 - the LUN ID of the LUN that stores the boot image
- to configure the LUN as the well known service name of your host.

Example

```
Topspin-360 (config) # fc srp initiator 00:05:ad:00:00:00:17:3c
00:00:00:00:00:00:00:00 bootup target 21:00:00:04:cf:f6:c2:ab lu
00:00:00:00:00:00:00:00
```

4. Boot your host.



NOTE: If the host does not boot successfully, reboot the host and enter interactive mode, then configure the boot order to begin with SRP boot or Well Known Service Name.

Use Element Manager

1. Remove or reconfigure any host that runs the existing image. Only one host on your network may boot a given image.
2. Physically connect your new host to your Server Switch.
3. Configure an ITL.



NOTE: To install an image with VFrame, do not configure ITLs manually.

- a. Open your host and record the GUID from the label on your HCA. Alternatively, you can boot your host and press the **x** key when the **Type 'x' to configure boot options** prompt appears during the boot process. The GUID of the host appears in the host display. Record the GUID and leave the host at this screen. You will refresh the screen later in this process.
- b. Launch Element Manager and open your Server Switch.

- c. From the **FibreChannel** menu, select **Storage Manager**. The **Storage Manager** window opens.
 - d. Click the **SRP Hosts** folder in the navigation tree.
 - e. Click the **Define New** button. The **Define New SRP Host** window opens.
 - f. Enter the GUID of your host (from [step a](#)) in the **Host GUID** field.
 - g. (Optional) Enter a text description for the host in the **Description** field.
 - h. Click the **Next >** button.
 - i. Click the **Finish** button. The new initiator appears under the **SRP Hosts** folder.
 - j. Click the new initiator under the **SRP Hosts** folder, then click the **Targets** tab.
 - k. Verify that the target that stores the image that you want to boot appears in the list.
 - l. Double-click the target in the list onto which you want to install the image. The **IT Properties** window opens.
 - m. Click the **...** button next to the **Port Mask** field. The **Select Ports** window opens.
 - n. Click the button of the port that connects the host to the storage, then click **OK**.
 - o. Click the **Apply** button in the **IT Properties** window, then close the window.
 - p. Click the **LUN Access** tab, then click the **Discover LUNs** button. The targets that your Server Switch sees appear in the **Available LUNs** column.
 - q. Click the target that stores the image that you want to boot, then click the **Add >** button.
 - r. Click the **Apply** button.
 - s. (Optional but recommended) Configure the Well Known Service Name.
 1. Return to the **Storage Manager** window of Element Manager and click the **General** tab.
 2. From the **Boot Target WWPN** field, select the boot target from the pulldown menu or manually enter it.
 3. From the **Boot FC LUN** field, select the boot LUN from the pulldown menu or manually enter it.
 4. Click the **Apply** button.
4. Boot your host.

BolB Prerequisites for Windows 2000

- Install software on your host (such as Paragon Partition Manager) that lets you copy your local drive to your remote FC storage.
- Install *only one* HCA in your host.
- Configure your SRP host as an initiator on your Server Switch (see *Topspin Fibre Channel Gateway User Guide*).
- (Optional) Configure your Server Switch to grant the host access to only one logical unit (LU) on the FC storage (see *Topspin Fibre Channel Gateway User Guide*). When you restrict access to one LU, you guarantee that you install the image to the correct LU.
- Configure the boot LUN with Element Manager (see *Topspin Fibre Channel Gateway User Guide*). You must configure the switch with the LU that you want the host to boot.
- Size your C drive so that it is the size of your LU or smaller, or verify that your LU is larger than your C drive.

Install an Image onto Fibre Channel Storage: Overview

To install an image onto FC storage and then boot that image, perform the following steps:

1. Log in to your Windows 2000 host, right-click the **My Computer** icon and select **Manage** from the right-click menu. The **Computer Management** window opens.
2. Click **Disk Management** and verify that your storage disk appears in the display.
3. Launch your partitioning software.
4. Verify that the C: drive is the same size as or smaller than the LU.
5. Copy the local disk to the remote disk.
6. Power down the server.
7. Remove the local drive from the host.
8. Change the boot order when you boot the host (refer to your hardware documentation for instructions).
9. Boot the server. The server boots from the SRP disk. The SRP disk becomes Drive C. Re-insert the local drive at any point thereafter.

Configuration Sample

The steps in this section detail how to install an image onto a storage LU and then configure a Windows 2000 server to boot that remote image. The instructions in this section assume the following:

- Only one HCA (“cougar” type, for this example) resides in the host.
- HCA drivers have been successfully installed on the host.
- The HCA runs BOIB firmware (see [“Check HCA Firmware Version” on page 71](#) and [“Upgrade Firmware” on page 25](#)).
- One IB cable connects the host to a Server Switch, and one Fibre Channel cable connects one port of a Fibre Channel gateway to one target port.
- The target storage device is a JBOD.
- The host is not configured as an initiator on the Server Switch.
- Element Manager is installed on the host and runs successfully on the Server Switch.

To perform this process, perform the following steps:

1. Launch Element Manager and open your Server Switch.
2. Click the **FibreChannel** menu and select **Storage Manager**. The **Storage Manager** window opens.
3. Click the **Storage** folder, then check the **Restricted** checkbox in the **LUN Access** field.
4. Uncheck the **Restricted** checkbox in the **Gateway Port Access** field, then click the **Apply** button.
5. Click the **SRP Hosts** folder, then click the **Define New** button. The **Define New SRP Host** window opens.
6. Select the GUID of your SRP host from the **Host GUID** field.
7. Enter a textual description in the **Description** field, click the **Next >** button, then click the **Finish** button. The entry for your host appears in the **SRP Hosts** folder.
8. Click the new initiator under the **SRP Hosts** folder, then click the **Targets** tab.
9. Verify that the WWPN of the target onto which you want to install the image appears in the list.
10. Double-click the WWPN (from [step i](#)) in the list of the target onto which you want to install the image. The **IT Properties** window opens.

11. Click the ... button next to the **Port Mask** field. The **Select Ports** window opens.
12. Click the button of the port that connects the host to the storage, then click **OK**.
13. Click the **Apply** button in the **IT Properties** window, then close the window.
14. Click the **LUN Access** tab, then click the **Discover LUNs** button. The targets that your Server Switch sees appear in the **Available LUNs** column.
15. Click the target (from [step i](#)) onto which you want to install an image, then click the **Add >** button.
16. Click the **Apply** button.
17. Click the **General** tab and select the WWPN of the boot target from the **Boot Target WWPN** pulldown menu.
18. Select the LUN ID of the boot LUN from the **Boot FC LUN** pulldown menu.
19. Log in to your Windows 2000 host, right-click the **My Computer** icon and select **Manage** from the right-click menu. The **Computer Management** window opens.
20. Click **Disk Management** and verify that your storage LU appears in the display.
21. Click **Start > Programs > Topspin Infiniband SDK > Utilities**. The **Utilities** window opens.
22. Enter `tvflash ../tavorFW/hca-cougar-a1-3-1-0-boot-srp.bin` at the command prompt to install the boot firmware on the HCA.



NOTE: The driver installation process places the boot firmware on your host when you install your HCA drivers.

23. Click **Start > Programs > Paragon Partition Manager > Partition Manager** to launch your partitioning software.

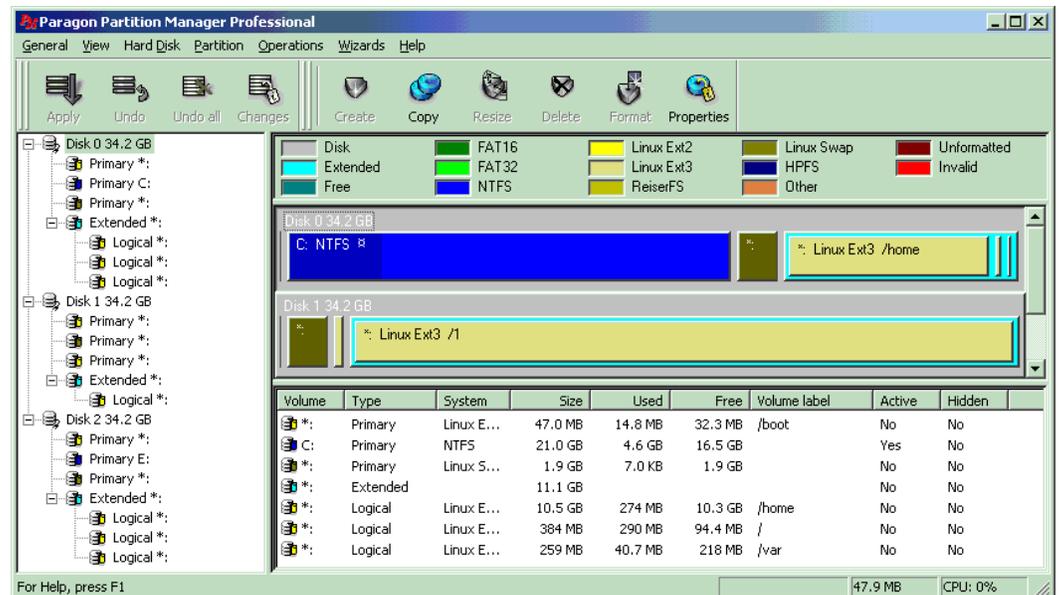


Figure 5-18: Partitioning Software

24. Verify that the C: drive is the same size as or smaller than the LU.

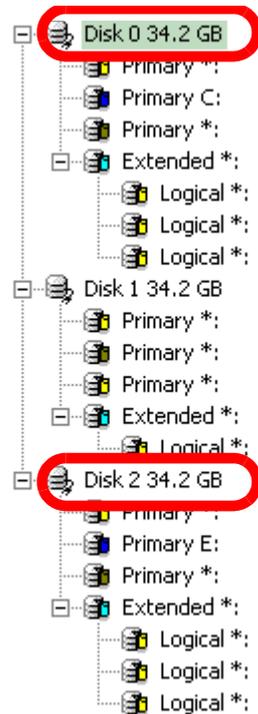
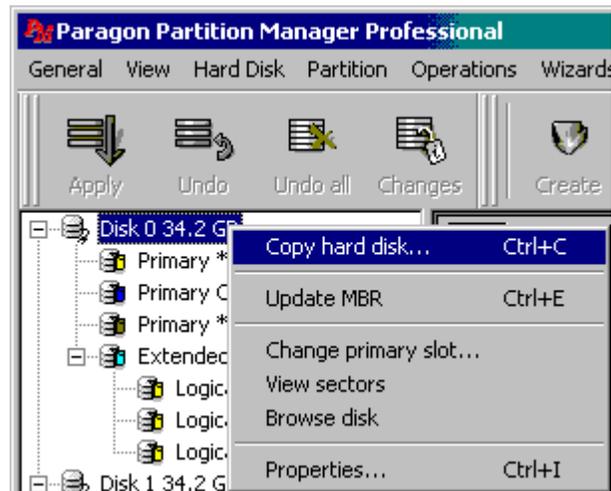


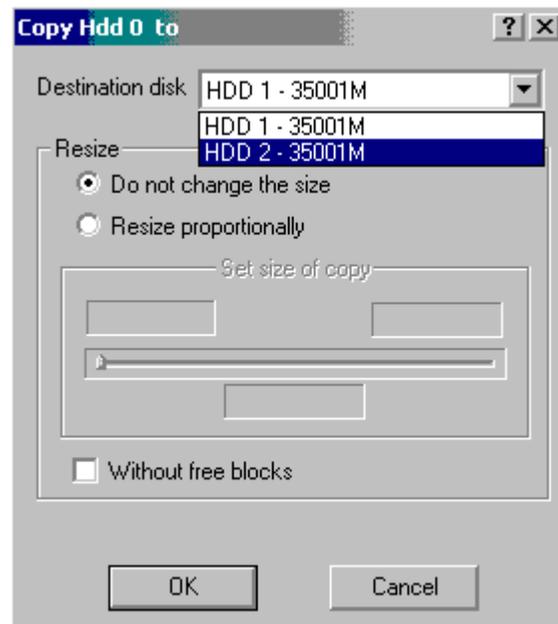
Figure 5-19: C: Drive (Disk 0) and LU (Disk 2) Sizes

25. Right-click the local disk (Disk 0) and select **Copy hard disk...** from the right-click menu. The **Copy Hdd 0** to window opens.

Example



26. Select your LU (Disk 2) from the **Destination disk** pulldown menu, then click **OK**.

Example

27. Reboot your server. Paragon software copies the local disk to the LU.
28. Power down the server.
29. Remove the local drive from the host.
30. Change the boot order when you boot the host (refer to your hardware documentation).
31. Boot the server. The server boots from the SRP disk. The SRP disk becomes Drive C. Re-insert the local drive at any point thereafter.

Booting a Host from PXE

The following sections appear in this chapter:

- [“About Booting Over PXE” on page 111](#)
- [“Configure Boot Over PXE” on page 112](#)

About Booting Over PXE

This document describes Topspin’s boot over InfiniBand using Pre-boot eXecution Environment (PXE). It contains high level overview of the boot process followed by a detailed step-by-step configuration procedure.

Sample PXE Hardware Setup

The setup used for the following example consists of one InfiniBand host booting Red Hat Linux Advanced Server release 2.1AS (Pensacola) and one server running Red Hat Enterprise Linux AS release 3 (Taroon) providing DHCP, TFTP and NFS services.

These services could be provided by a single or separate network nodes connected either to the InfiniBand or Ethernet network. This document covers both cases. If some of the services are provided by Ethernet node a Topspin Ethernet Gateway is required. A sample configuration for the Ethernet Gateway is provided as well.

Overview of the PXE Process

The host BIOS reads the configuration from non-volatile memory to determine the boot device order. In order to boot over InfiniBand using PXE the first entry for the boot order should be PXE, by default. Use the `tvflash -o` command to set the boot order without entering the BIOS, as described in [“Configure the Booting Host” on page 113](#).

The host BIOS sends a DHCP request to a PXE-enabled DHCP server to obtain initial configuration and boot parameters. This includes IP address and boot loader file name.

The host BIOS loads the boot loader from TFTP server and starts it.

Boot loader reads its configuration file from TFTP server and parses it to get boot options, Linux kernel and initrd file names.

The boot loader loads the Linux kernel and executes it, passing some options from the configuration file.

The Linux kernel performs the initialization procedures and loads initrd image from TFTP server as a temporary root file system before mounting the root file system.

The initrd image startup script starts and loads Topspin's InfiniBand stack and IPoIB driver. At the end it changes root file system to a NFS volume.

The Linux initialization continues with the NFS root file system, using the IPoIB interface until it is up and running.

Configure Boot Over PXE

To perform a PXE boot with the Boot over IB feature (BoIB-PXE), you must perform the following steps:

[“Verify Prerequisites” on page 112](#)

[“Perform Hardware Setup” on page 112](#)

[“Verify PXE-Enabled Firmware for Booting Server” on page 113](#)

[“Configure the Booting Host” on page 113](#)

[“Configure the DHCP Server” on page 114](#)

[“Configure the TFTP Server” on page 115](#)

[“Configure Boot Loader” on page 116](#)

[“Configure Topspin Ethernet Gateway” on page 117](#)

Verify Prerequisites

The following requirements must be met to perform boot over PXE:

- a host with a Topspin host channel adapter (HCA)
- a server with a PXE-enabled Topspin host channel adapter (HCA). The booting server provides DHCP, TFTP and NFS services. These services could be provided by a single or separate network nodes connected either to the InfiniBand or Ethernet network. This document covers both cases. If some of the services are provided by Ethernet node a Topspin Ethernet Gateway is required.
- a Server Switch, such as a Topspin 90 or Topspin 360, with an optional Ethernet gateway

Perform Hardware Setup

1. Install the Topspin HCAs into your host and server.



NOTE: Only Topspin “Cougar”, “Cougar Cub” and “Lion Cub” family HCAs support the Boot over IB feature. To verify that your HCA supports BoIB, run the `tvflash -i` command and verify that Cougar or Lion Cub appears in the output.

2. Install and connect all required hardware.
3. Connect the HCA to an InfiniBand port on your Server Switch.

Verify PXE-Enabled Firmware for Booting Server

4. View the firmware on the HCA of the booting server and load the correct firmware onto the HCA, if needed. Your firmware must be PXE-enabled.

You can upgrade standard HCAs to run Boot over IB-capable and PXE-capable firmware. To upgrade the firmware on your HCA, refer to the HCA/Host -Side Driver User Guide for detailed information.

- a. View your firmware on the booting server using the `tvflash -i` command. Your primary image should be have a `.PXE` extension.

Example

```
# /usr/local/topspin/sbin/tvflash -i
READ_VPD: timed out at off 0
HCA #0: Found MT23108, Cougar, revision A1 (firmware autoupgrade)
  Primary image is v3.02.0000 build 2.5.0.204, with label
'HCA.Cougar.A1.Boot.PXE'
Secondary image is v3.01.0000, with label 'HCA.Cougar.A1.Boot'
```

- b. Upgrade your firmware if it is not PXE-capable.

Example

```
# /usr/local/topspin/sbin/tvflash <firmware_name>
New Node   GUID = 0005ad0000001702
New Port1  GUID = 0005ad0000001703
New Port2  GUID = 0005ad0000001704
Programming Tavor Microcode... Flash Image Size = 408452
Failsafe   100%
Erasing    100%
Writing    100%
Verifying  100%
Flash verify passed!
```

Configure the Booting Host

The host drivers provide a feature (`tvflash -o`) that allows you to change the default boot option without using the boot ROM itself. If you decide to set the default option without using the automatic `tvflash -o` option, you can do this manually. The procedure is BIOS-dependent. Set the boot device order so that boot over PXE is the one that the BIOS selects to boot from. This step requires entering motherboard setup program, finding the menu item specifying the boot order and changing it accordingly, if needed.

5. View the current default boot option by entering the `tvflash -o` command.

Example

```
[root@ home]# /usr/local/topspin/sbin/tvflash -o
auto_upgrade=yes
boot_enable_port_1=yes
boot_enable_port_2=yes
boot_service_scan=yes
boot_type=well_known
boot_saved_port=0
boot_saved_ioc_num=0
boot_saved_dgid=0000:0000:0000:0000:0000:0000:0000:0000
boot_saved_service_name=0000:0000:0000:0000
```

6. Look for the `boot_type` field. In this example, the `boot_type` is set to “well_known,” so the boot hosts will try the well-known service name option by default.
7. Change the default boot option to PXE, if the default `boot_type` is not PXE. Use the `tvflash -o boot_type=pxe` command.

Setting the value to pxe requires the use of a boot ROM which is PXE capable.

Example

```
[root@ home]# /usr/local/topspin/sbin/tvflash -o boot_type=pxe
Writing [=====]
[root@ home]# /usr/local/topspin/sbin/tvflash -o
auto_upgrade=yes
boot_enable_port_1=yes
boot_enable_port_2=yes
boot_service_scan=yes
boot_type=pxe
boot_saved_port=0
boot_saved_ioc_num=0
boot_saved_dgid=0000:0000:0000:0000:0000:0000:0000:0000
boot_saved_service_name=0000:0000:0000:0000
```

Configure the DHCP Server

8. Verify that the DHCP server supports PXE extensions.

The following example shows how to configure Linux DHCP server using `dhcp` package.

- a. The DHCP server configuration `/etc/dhcpd.conf` file must contain a section describing the IP subnet of the booting host. The minimum required information is as follows:

Example

```
ddns-update-style none;

subnet 40.0.0.0 netmask 255.255.0.0 {
    next-server 40.0.0.101;
    filename "/pxelinux.0";

    range dynamic-bootp 40.0.0.50 40.0.0.100;
    always-broadcast on;
}
```

The important options are next-server, filename and always-broadcast.

- The Next-server and Filename options specify the TFTP server and file name of the initial boot loader.
 - The always-broadcast option should be turned on. Turning this option on causes the DHCP server to broadcast DHCP OFFER reply, rather than unicast. This is part of DHCP over IB specification, and is required for the IB DHCP client to acquire the IP address. IP addresses and ranges can be changed to match target setup. Note that the broadcast option must be specified even if the DHCP server is on the Ethernet network.
- b. (Optional) To limit the DHCP server to specific IP interfaces, include the following statement in the `/etc/sysconfig/dhcpd` file:

DHCPDARGS=<interface>

In the example below, the DHCP server will be active only on the `ib0` interface.

Example

```
[root@ home]# /etc/sysconfig/dhcpd DHCPDARGS="ib0 lo0"
```

- c. Activate the DHCP server after the configuration files are prepared. To start the server, use the `service dhcpd start` command as root.

Example

```
[root@ home]# # service dhcpd start
Starting dhcpd:
[ OK ]
```

- d. Make the DHCP configuration persistent across reboots by using the `chkconfig` command.

Example

```
# chkconfig --list dhcpd
dhcpd          0:off  1:off  2:on   3:on   4:on   5:on   6:off
```

- e. Verify that the DHCP server is running.

Example

```
# ps -ef | grep dhcp
root      4069    1  0 12:06 ?        00:00:00 /usr/sbin/dhcpd ib0 lo0
```

Configure the TFTP Server

The following steps assume you are using the Linux `tftp-server` package as a TFTP server.

In this example, the setup `tftp` service is handled by `xinetd` daemon. The corresponding `xinetd` configuration file must be edited in order to enable the service. The following is a sample `/etc/xinetd.d/tftp` file.

9. Enable TFTP service. Edit the xinetd.d file to enable TFTP service with root directory /tftpboot. All file paths have to be relative to this directory.

Example

```
service tftp
{
    socket_type          = dgram
    protocol            = udp
    wait                = yes
    user                 = root
    server               = /usr/sbin/in.tftpd
    server_args          = -s /tftpboot
    disable              = no
    per_source           = 11
    cps                  = 100 2
    flags                = IPv4
}
```

10. Activate the configuration by restarting the inetd daemon.

Example

```
# service xinetd restart
Stopping xinetd:           [ OK ]
Starting xinetd:          [ OK ]
```

TFTP service is now active and the initial boot loader, Linux kernel and initrd image can be placed under TFTP root directory.

Configure Boot Loader

Note: The boot loader used in this example is from syslinux distribution <http://syslinux.zytor.com>.

Only the PXE boot loader is used.

The file name is *pxelinux.0*.

11. Copy the boot loader file to the tftpboot directory stated in the DHCP configuration file. The boot loader expects to find its configuration file on the TFTP server under pxelinux.cfg directory.
12. Create /tftpboot/pxelinux.cfg directory and create a configuration file.

The boot loader configuration could be for a specific host, for a group of hosts, or for all hosts.

Creating Individual Configuration Files

To create individual configuration files, the file name must match the corresponding host IP address as a 32 bit hex number. For example for a host with IP address 40.0.0.99, the file name would have to be 28000063.

Creating Configuration Files for All Hosts

To create a single configuration file to be used for all hosts (as in the example below), the file name must be *default*.

a. The sample below file configures boot loader to:

load linux kernel with file name linux and initrd image file name initrd from the /tftpboot directory on the TFTP server. It also passes some kernel parameters to the Linux kernel.

Example

```
serial 0 38400
default rhel3
prompt 1
timeout 100

label rhel3
kernel linux
append root=/dev/cciss/c0d0p3 initrd=initrd console=ttyS0,38400
```

Configure Topspin Ethernet Gateway

If you have configured your services so that DHCP, TFTP or NFS servers are on the Ethernet network, a Topspin Ethernet Gateway must to be configured to bridge the IP subnet.

The following sample configuration is for an Ethernet Gateway located in a Topspin 360 chassis:

- The Ethernet Gateway is installed in slot 6.
- One bridging group is configured between Ethernet port 5 on the gateway and the second InfiniBand port, which is internal.
- No VLAN tagging is configured and the default partition key (P_KEY) is used on the InfiniBand side.
- Note that broadcast forwarding must be enabled in order for the DHCP server to work across the gateway.

Example

```
Topspin-360# configure
Topspin-360(config)# bridge-group 2 subnet-prefix 40.0.0.0 16
Topspin-360(config)# bridge-group 2 broadcast-forwarding
Topspin-360(config)# interface ethernet 6/5
Topspin-360(config)# bridge-group 2
Topspin-360(config)# interface gateway 6/2
Topspin-360(config-if-gw-2/2)# bridge-group 2
Topspin-360(config-if-gw-2/2)# exit
```


Specifications and Compliance

This chapter details the HCA Expansion Card specifications and compliance certifications.

Physical and Environmental Specifications

[Table 7-1](#) lists temperature and cooling conditions.

Table 7-1: Temperature and Cooling

Condition	Range
Operating Temperature	5 to 67 degrees Celsius, local ambient temperature
System Non-Operating	-40 to 90 degrees Celsius
Operating Altitude	0 to 10,000 ft
Non-Operating Altitude	0 to 50,000 ft

[Table 7-2](#) lists temperature and cooling conditions.

Table 7-2: Humidity

Condition	Range
Operating Humidity	8 to 80% non-condensing
Operating Gradient	Max per 60 min: 10%
Non-Operating Humidity	5 to 80%, non-condensing

Electromagnetic Compatibility

The HCA Expansion Card meets the following EMC “Class A” standards, when installed in and operated within a closed chassis:

- USA: CFR 47 Part 15, Subpart B
- Canada: ICES-003 Issue 2
- Europe: EN 61000-3-2 (Harmonics), EN 61000-3-3 (Flicker), EN 55022 1998, EN 55024 1994
- Japan: VCCI-V3/97.04

Index

B

BoIB-SAN overview	69
boot order	
set for boot over SAN	85
boot over SAN	
set boot order	85

C

CLI, using to place image on storage	73
configuring IPoIB	35
configuring MPI	47
configuring SDP	50
Configuring SRP	45
Continue Anyway button	23

D

del	45
DISABLED	62
disabling interfaces	57
driver installation	17

E

element manager, using to place image on storage	72
enabling interfaces	57

F

features	2
firmware installation	17
firmware, updated	71
firmware, upgrading	25

G

GID	62
global access, restricting	72, 73

H

hardware installation	5
hardware installation, FC	70
hardware Setup, FC	70
HCA expansion cards, about	2
HCA firmware version	71

I

IB switch module	1
image, placing on storage with CLI	73
image, placing on storage with element manager	72
installing HCA drivers	17
installing HCA firmware	17
installing HCA software	17
installing the HCA hardware	5
interfaces, enabling and disabling	57
introduction	1
IPoIB protocol	3
IPoIB, configuring	35
ipoibcfg	45
ITL, configuring	72, 74

L

LID	
view port	62
view subnet manager	62
LINK_UP	62

M	
Message Passing Interface	4
Microsoft HCT suite	22
MPI	4
MPI, configuring	47
O	
overview of BoIB-SAN	69
P	
phy state	
view port logical link	62
PKEY	
view partitions	62, 63
port	
GID	62
view logical link state	62
view partitions	62, 63
port receive data	67
port receive packets	67
port state	
change	62, 66
view	62, 66
port transmit packets	67
properties	
view HCA properties	60, 64
protocols	
IPoIB	3
MPI	4
SDP	3
SRP	3
uDAPL	3
protocols, supported	3
R	
restricting global access	72, 73
S	
safety information	xi
SCSI RDMA Protocol	3
SDP protocol	3
SDP, configuring	50
server switch module	1
SM_LID	62
sockets direct protocol	3
software installation	17
SRP	3
SRP, configuring	45
subnet manager	
view LID	62
supported protocols	3
switch module	1
T	
topology	
dual-switch	2
single-switch	2
tvflash -o	
set boot order for SAN	85
U	
uDAPL	3
upgrading firmware	25
User Direct Access Programming Library	3
utilities	45
V	
verifying driver installation	20
version, HCA firmware	71
vstat	55
W	
well known service name, configuring ...	73, 78, 106
WHQL testing	22
Windows signature	22