



Release Notes for Cisco OFED Release 1.1

Release Date: December 13, 2006
Part Number: OL-11537-01

Contents

This document contains these sections:

- [Contents, page 1](#)
- [Introduction, page 1](#)
- [System Requirements, page 2](#)
- [New and Changed Information, page 4](#)
- [Caveats, page 6](#)
- [Related Documentation, page 7](#)
- [Obtaining Documentation, page 9](#)
- [Documentation Feedback, page 10](#)
- [Cisco Product Security Overview, page 10](#)
- [Obtaining Technical Assistance, page 11](#)
- [Obtaining Additional Publications and Information, page 13](#)

Introduction

These release notes describe the features and known issues for the Cisco OpenFabrics Enterprise Distribution (OFED) InfiniBand Host Drivers 1.1 for Linux.



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2006 Cisco Systems, Inc. All rights reserved.

System Requirements

This section describes the system requirements for this software release.

Cisco supports OFED on all existing Cisco InfiniBand HCAs, Cisco InfiniBand switches, and Cisco Ethernet and Fibre Channel Gateways.

Determining the Software Version

If InfiniBand drivers are already installed on the host, they may be installed in one of several locations.

To determine the version of the Cisco InfiniBand host drivers, log in to the host and enter the following commands at the shell prompt. If the first command produces output, the Cisco Commercial InfiniBand host drivers are installed. If the second command produces a version number, OFED host drivers are installed.

```
host$ rpm -qa | grep topspin
topspin-ib-mpi-rhel4-3.2.0-118
topspin-ib-mod-rhel4-2.6.9-34.ELsmp-3.2.0-118
topspin-ib-rhel4-3.2.0-118
host$ grep OFED /usr/local/ofed/BUILD_ID
OFED-1.0
```

Upgrading to a New Software Release

To verify that you are running the latest available release, compare your version against the latest version on the Cisco support website at <http://www.cisco.com/cgi-bin/tablebuild.pl/sfs-linux>. After registering your product, you should have received a username and password that grant you access to this site.

Switch software and Linux host drivers are released and packaged separately. Cisco OFED 1.1 requires that all switches first be upgraded to TopspinOS 2.4.0 or higher. Switches should be upgraded before the InfiniBand hosts.

Use the **ofedinstall** command to install Cisco OFED 1.1 on your host. There can only be one set of InfiniBand drivers installed on a host at any time. The **ofedinstall** command will first uninstall any Cisco Commercial InfiniBand host drivers or other OpenFabrics host drivers. If the host has other non-OpenFabrics drivers, those must be uninstalled first before running **ofedinstall**.

To enable userspace components such as MPI, **ofedinstall** adds the following entries to `/etc/security/limits.conf`. You may tune the value unlimited to a specific amount of RAM if desired.

```
* soft memlock unlimited
* hard memlock unlimited
```

For general information about upgrading to a new software release, see the Installing Host Drivers chapter in the *Cisco OpenFabrics Enterprise Distribution InfiniBand Host Drivers User Guide for Linux*.

Supported Kernels

Cisco OFED 1.1 InfiniBand host drivers are supported on the following kernels:

- Red Hat Enterprise Linux 4 (RHEL4)
 - 2.6.9-34.ELsmp (Update 3) for i686
 - 2.6.9-34.ELhugemem (Update 3) for i686

- 2.6.9-34.ELsmp (Update 3) for x86_64
- 2.6.9-34.ELlargesmp (Update 3) for x86_64
- 2.6.9-34.EL (Update 3) for ia64
- 2.6.9-34.ELlargesmp (Update 3) for ia64
- 2.6.9-34.EL (Update 3) for ppc64
- 2.6.9-34.ELlargesmp (Update 3) for ppc64
- 2.6.9-42.ELsmp (Update 4) for i686
- 2.6.9-42.ELhugemem (Update 4) for i686
- 2.6.9-42.ELsmp (Update 4) for x86_64
- 2.6.9-42.ELlargesmp (Update 4) for x86_64
- 2.6.9-42.EL (Update 4) for ia64
- 2.6.9-42.ELlargesmp (Update 4) for ia64
- 2.6.9-42.EL (Update 4) for ppc64
- 2.6.9-42.ELlargesmp (Update 4) for ppc64
- SUSE Linux Enterprise Server (SLES10)
 - 2.6.16.21-0.8-smp for i686
 - 2.6.16.21-0.8-smp for x86_64
 - 2.6.16.21-0.8-default for ia64
 - 2.6.16.21-0.8-ppc64 for ppc64

If you are using a different kernel, you will need to compile OFED from source code. To compile OFED source code, unpack the OFED tarball in `src/` in the Cisco OFED ISO image and follow the instructions in `README.txt`. If you compile OFED from source code, Cisco will support you on a best-effort basis, but we recommend that you use the pre-built binary RPMs if possible. More information about OFED is available in `docs/` in the ISO image, `/usr/local/ofed/docs` on a host with the drivers installed, <http://www.openfabrics.org>, and the OpenFabrics mailing list `openib-general` and its mail archives.

Supported MPI Implementations

Cisco OFED 1.1 InfiniBand host drivers have been tested with the following MPI implementations:

- OFED MVAPICH 0.9.7 (included with Cisco OFED 1.1)
- Open MPI 1.1.2 (included with Cisco OFED 1.1)
- Intel MPI 2.0.1 refresh
- HP MPI 2.2.5

Other MPI implementations support OFED, for more information contact the MPI vendor.

Supported Compilers for MPI

Cisco OFED 1.1 InfiniBand host drivers are supported with the following compilers:

- GNU C, C++, Fortran 77, and Fortran 90
- Intel 9.1 C, C++, Fortran 77, and Fortran 90

Supported HCA Firmware

Cisco OFED 1.1 InfiniBand host drivers include the following InfiniBand HCA firmware, which is based on the Mellanox firmware from http://www.mellanox.com/support/firmware_table.php:

HCA Type	Description	Firmware Version
Cougar	2 port PCI-X (not low-profile)	3.5.000
Cougar Cub	2 port PCI-X	3.5.000
Lion Cub	2 port PCIe	4.7.600
Lion Mini	2 port PCIe memfree	5.1.400
Cheetah	1 port PCIe memfree	1.1.000
Cheetah DDR	1 port PCIe memfree DDR	1.1.000
IBM BladeCenter PCI-X	uses Cougar firmware	see Cougar
IBM BladeCenter PCIe	based on Lion Cub	4.7.600
Dell PowerEdge 1X55	based on Lion Cub	4.7.600

New and Changed Information

Cisco OFED 1.1 is a major release that introduces significant new features and documentation. It includes the following changes:

- Support for RHEL4 Update 4
- Support for SLES10
- Support for 20 Gbps InfiniBand Double Data Rate HCAs
- Improved SDP performance and stability
- Many bug fixes

Changes From Release 1.0 to 1.1

This section describes the new features and resolved caveats since the Cisco OFED 1.0 release.



Note

The ID number from the OpenFabrics Defect Tracking System, if applicable. The current status of all issues is available online at <http://www.openfabrics.com/bugzilla>.

General

- 217
Support has been added for 20 Gbps InfiniBand Double Data Rate HCAs. Note that the longest DDR copper IB cable supported at this time is 8 meters.
- 43, 101
Linux man pages have been added for most utility and diagnostic programs.
- OFED 1.1 userspace programs are binary compatible with OFED 1.0.

- RDS has been removed from OFED 1.1.

IPoIB

- IPoIB throughput, latency, and CPU utilization are improved with OFED compared to the Cisco Commercial Linux host drivers.
- Support for IPoIB High Availability was added, but is not yet documented or supported by Cisco.

MPI

- Open MPI in Cisco OFED has been upgraded from 1.1 to 1.1.2.
- 173
The OSU MPI benchmarks have been upgraded to the newer `osu_latency`, `osu_bw`, `osu_bibw`, and `osu_bcast` programs.
- The Pallas MPI benchmarks have been upgraded to the newer Intel MPI benchmarks.
- MVAPICH small message throughput has been improved.
- 80
Open MPI is now supported on ppc64.
- 135
MVAPICH is now supported on ppc64.
- 191
MVAPICH now supports gfortran GNU Fortran 90.
- 267
On SLES10, **ofedinstall** will comment out an entry in `/etc/hosts` for 127.0.0.2, otherwise MVAPICH will not function.

SDP

- 24, 109
SDP performance and stability are now improved. SDP latency is improved with OFED compared to the Cisco Commercial Linux host drivers.

SRP

- Support for SRP High Availability was added, but is not yet documented or supported by Cisco.

uDAPL

- uDAPL scalability has been improved for use by Intel MPI.

Caveats

This section describes temporary limitations of this release. These restrictions will be resolved in a future release of this product.

General

- Cisco does not support the OpenSM InfiniBand subnet manager.
- 270
tvflash does not work with HCA recovery jumper.
- 183
The IB node description will sometimes not have the hostname in it.
- 292
tvflash does not work on SLES10 ia64.
- 265
The OFED uninstall script `uninstall.sh` does not work if multiple `kernel-ib` packages are installed. The workaround is to uninstall the `kernel-ib` packages manually before running `uninstall.sh`.

IPoIB

- 263
Support for IPoIB High Availability was added, but is not yet documented or supported by Cisco due to stability problems.
- IPoIB performance varies greatly by type of motherboard. The best performance will be achieved with the latest Intel and AMD processors and PCIe HCAs.
- 266
IPoIB IP multicast is not supported on RHEL4 U4. U4 IP multicast can only receive from U4 senders, and U4 senders sent traffic to all nodes, not just nodes that have joined the multicast address. U4 multicast traffic can cause a Cisco InfiniBand switch to reboot.
- RHEL4 `tcpdump` command does not work with OFED IPoIB.
- RHEL4 and SLES10 `ifconfig -a` commands do not display the IPoIB interface address correctly. The workaround is to use the `ip link show ib0` command instead.
- 269
IPoIB does not interoperate with the `mgmt-ib` IPoIB management interface on the Cisco SFS 3001 and Topspin 90 switch.

MPI

- 188
Open MPI performance on some Pallas/Intel MPI collective benchmarks is poor.
- 187
Open MPI does not fully support MPI-2 Remote Memory Access.

- 74
Open MPI latency is poor on ia64.
- 258
MVAPICH does not support GNU Fortran 90 on ppc64.

SDP

- 137
OFED SDP does not interoperate with Cisco Commercial Linux InfiniBand Host Driver SDP.
- 108
SDP throughput with message sizes greater than 64 KB may decrease as the message size increases.
- SDP performance varies greatly by type of motherboard. The best performance will be achieved with the latest Intel and AMD processors and PCIe HCAs.
- 38
OFED SDP does not support Linux AIO.
- 25
There is no `#define` for `AF_INET_SDP`, code that uses SDP without `libsdp` must manually add `#define AF_INET_SDP 27` at the beginning of a program.
- 294
SDP does not support using the `connect()` function with `AF_INET_SDP`. The workaround is to use `AF_INET` for this function instead. `AF_INET_SDP` must be used with the `socket()` and `bind()` functions.

SRP

- Support for SRP High Availability was added, but is not yet documented or supported by Cisco.

uDAPL

- 48
uDAPL is not available on ppc64.

Related Documentation

The following list describes the documentation available with TopspinOS 2.9.0, which is available in electronic form and printed form upon request.



Note

Documentation is included on the TopspinOS 2.9.0 Server Switch CD Image.

You can download the latest documentation updates on the Cisco support site at http://www.cisco.com/en/US/partner/products/ps6418/tsd_products_support_category_home.html.

- *InfiniBand Hardware Installation and Cabling Guide*

- *Cisco SFS 7000P Switch Installation and Configuration Note*
- *Cisco SFS 7000P Hardware Installation Guide*
- *Cisco SFS 7008P Switch Installation and Configuration Note*
- *Cisco SFS 7008P Hardware Installation Guide*
- *Release Notes for TopspinOS Release 2.9.0*
- *Cisco SFS 7000 Series Product Family Chassis Manager User Guide*
- *Cisco SFS 7000 Series Product Family Element Manager User Guide*
- *Cisco SFS 7000 Series Product Family Command Reference Guide*

Service and Support

For additional support, you must first register your product at <http://www.cisco.com>. After registering, you may contact your supplier for support, or Cisco directly.

Refer to the [“Obtaining Technical Assistance” section on page 11](#) in this document.

When you call Cisco Technical Support or use the Cisco Technical Support website at <http://www.cisco.com>, be prepared to provide the following information to support personnel:

General Information

- Technical Support registration number, if applicable
- Error messages received
- Detailed description of the problem and specific questions
- Description of any troubleshooting steps already performed and results

Server Configuration

- Type of server, chip set, CPU, amount of RAM, and number of nodes
- Attached storage devices (output from `cat /proc/scsi/scsi`)
- InfiniBand configuration (output from `hca_self_test`)

Chassis Configuration

- Chassis model
- Output from the `show running-status all` command

Chassis Serial Number

The chassis serial number and corresponding bar code are provided on the serial number label, as shown in this example:

Model: TS360



SN UST323XXXXXXXXXX

This chassis serial number can be found on the bottom of the chassis or the outside of the chassis box packaging. It can also be found in the output of the **show backplane** command.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. This section explains the product documentation resources that Cisco offers.

Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/techsupport>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

The Product Documentation DVD is a library of technical product documentation on a portable medium. The DVD enables you to access installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the HTML documentation and some of the PDF files found on the Cisco website at this URL:

<http://www.cisco.com/univercd/home/home.htm>

The Product Documentation DVD is created and released regularly. DVDs are available singly or by subscription. Registered Cisco.com users can order a Product Documentation DVD (product number DOC-DOCDVD= or DOC-DOCDVD=SUB) from Cisco Marketplace at the Product Documentation Store at this URL:

<http://www.cisco.com/go/marketplace/docstore>

Ordering Documentation

You must be a registered Cisco.com user to access Cisco Marketplace. Registered users may order Cisco documentation at the Product Documentation Store at this URL:

<http://www.cisco.com/go/marketplace/docstore>

If you do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Documentation Feedback

You can provide feedback about Cisco technical documentation on the Cisco Technical Support & Documentation site area by entering your comments in the feedback form available in every online document.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you will find information about how to do the following:

- Report security vulnerabilities in Cisco products
- Obtain assistance with security incidents that involve Cisco products
- Register to receive security information from Cisco

A current list of security advisories, security notices, and security responses for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you have identified a vulnerability in a Cisco product, contact PSIRT:

- For emergencies only — security-alert@cisco.com

An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- For nonemergencies — psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

**Tip**

We encourage you to use Pretty Good Privacy (PGP) or a compatible product (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.x through 9.x.

Never use a revoked encryption key or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

If you do not have or use PGP, contact PSIRT to find other means of encrypting the data before sending any sensitive material.

Product Alerts and Field Notices

Modifications to or updates about Cisco products are announced in Cisco Product Alerts and Cisco Field Notices. You can receive Cisco Product Alerts and Cisco Field Notices by using the Product Alert Tool on Cisco.com. This tool enables you to create a profile and choose those products for which you want to receive information.

To access the Product Alert Tool, you must be a registered Cisco.com user. (To register as a Cisco.com user, go to this URL: <http://tools.cisco.com/RPF/register/register.do>) Registered users can access the tool at this URL: <http://tools.cisco.com/Support/PAT/do/ViewMyProfiles.do?local=en>

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support & Documentation Website

The Cisco Technical Support & Documentation website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support & Documentation website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

**Note**

Use the **Cisco Product Identification Tool** to locate your product serial number before submitting a request for service online or by phone. You can access this tool from the Cisco Technical Support & Documentation website by clicking the **Tools & Resources** link, clicking the **All Tools (A-Z)** tab, and then choosing **Cisco Product Identification Tool** from the alphabetical list. This tool offers three search options: by product ID or model name; by tree view; or, for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

**Tip**

Displaying and Searching on Cisco.com

If you suspect that the browser is not refreshing a web page, force the browser to update the web page by holding down the Ctrl key while pressing F5.

To find technical information, narrow your search to look in technical documentation, not the entire Cisco.com website. On the Cisco.com home page, click the **Advanced Search** link under the Search box and then click the **Technical Support & Documentation** radio button.

To provide feedback about the Cisco.com website or a particular technical document, click **Contacts & Feedback** at the top of any Cisco.com web page.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests, or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411

Australia: 1 800 805 227

EMEA: +32 2 704 55 55

USA: 1 800 553 2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—An existing network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of the network is impaired while most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Online Subscription Center is the website where you can sign up for a variety of Cisco e-mail newsletters and other communications. Create a profile and then select the subscriptions that you would like to receive. To visit the Cisco Online Subscription Center, go to this URL:
<http://www.cisco.com/offer/subscribe>
- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco channel product offerings. To order and find out more about the *Cisco Product Quick Reference Guide*, go to this URL:
<http://www.cisco.com/go/guide>
- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:
<http://www.cisco.com/go/marketplace/>
- Cisco Press publishes a wide range of general networking, training, and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:
<http://www.ciscopress.com>
- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the *Internet Protocol Journal* at this URL:
<http://www.cisco.com/ipj>
- Networking products offered by Cisco Systems, as well as customer support services, can be obtained at this URL:
<http://www.cisco.com/en/US/products/index.html>

- Networking Professionals Connection is an interactive website where networking professionals share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:
<http://www.cisco.com/discuss/networking>
- “What’s New in Cisco Documentation” is an online publication that provides information about the latest documentation releases for Cisco products. Updated monthly, this online publication is organized by product category to direct you quickly to the documentation for your products. You can view the latest release of “What’s New in Cisco Documentation” at this URL:
<http://www.cisco.com/univercd/cc/td/doc/abtunibd/136957.htm>
- World-class networking training is available from Cisco. You can view current offerings at this URL:
<http://www.cisco.com/en/US/learning/index.html>

This document is to be used in conjunction with the documents listed in the “[Related Documentation](#)” section.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2006 Cisco Systems, Inc. All rights reserved.