

2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter

IBM BladeCenter at-a-glance guide

The 2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter is a dual port InfiniBand Host Channel Adapter (HCA) based on proven Mellanox ConnectX IB technology. This HCA, when combined with the QDR switch, delivers end-to-end 40 Gb bandwidth per port. This solution is ideal for low latency, high bandwidth, performance-driven server and storage clustering applications in a High Performance Compute environment. The adapter uses the CFFh form factor and can be combined with a CIOv or CFFv adapter to get additional SAS, Fibre Channel, or Ethernet ports.

Figure 1 shows the expansion card.



Figure 1. 2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter

Did you know?

InfiniBand is a scalable high performance fabric that was used for Petascale computing. Roadrunner is the largest supercomputer of the world, breaking the barrier of 1000 trillion operations per second. Roadrunner is based on Mellanox ConnectX DDR adapters. QDR is the next generation, which offers twice the bandwidth per port.

Part number information

Table 1 shows the part numbers to order the 2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter.

Table 1. Part number and feature code for ordering

Description	Part number	Feature codes
2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter	46M6001	0056

The part number includes the following items:

- One 2-Port 40 Gb InfiniBand Expansion Card (CFFh)
- Documentation CD-ROM
- Important Notices flyer

Features and specifications

The 2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter includes the following features and specifications:

- The 2-Port 40 Gb InfiniBand Expansion Card features include:
 - Form-factor: CFFh
 - Host interface: PCI-E x8 Gen 2 (5.0GT/s): 40+40 Gbps bidirectional bandwidth
 - Dual 4X InfiniBand ports at speeds 10 Gbps, 20 Gbps, or 40 Gbps per port
 - 6.5 GBps bidirectional performance
 - RDMA, Send/Receive semantics
 - Hardware-based congestion control
 - Atomic operations
 - 16 million I/O channels
 - 256 to 4 Kb MTU
 - 1 GB messages
 - 9 virtual lanes: 8 data + 1 management
 - 1us MPI ping latency
 - CPU offload of transport operations
 - End-to-end QoS and congestion control
 - TCP/UDP/IP stateless offload
- Enhanced InfiniBand features
 - Hardware-based reliable transport
 - Hardware-based reliable multicast
 - Extended Reliable Connected transport
 - Enhanced Atomic operations
 - Fine grained end-to-end QoS
- Hardware-based I/O virtualization features
 - Single Root IOV
 - Address translation and protection
 - Multiple queues per virtual machine
 - VMware NetQueue support

- Protocol support
 - Open MPI, OSU MVAPICH, HP MPI, Intel MPI, MS MPI, Scali MPI
 - IPoIB, SDP, RDS
 - SRP, iSER, FCoIB and NFS RDMA

Operating environment

2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter supports the following environment:

- Temperature
 - 10 to 52 °C (50 to 125.6 F) at an altitude of 0 to 914 m (0 to 3,000 ft)
 - 10 to 49 °C (50 to 120.2 F) at an altitude of 0 to 3000 m (0 to 10,000 ft)
- Relative humidity
 - 8% to 80% (noncondensing)

Supported servers and I/O modules

Table 2 lists the IBM BladeCenter servers that the 2-Port 40 Gb InfiniBand Expansion Card for IBM BladeCenter supports.

Table 2. Supported servers

Expansion card	HS12	HS21	HS21 XM	HS22	HS22V	HX5	LS21	LS22	LS41	LS42	JS12	JS21	JS22	JS23/JS43	PS700/1/2
2-Port 40 Gb InfiniBand Expansion Card (CFFh)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N

Figure 2 shows where the CFFh card is installed in a BladeCenter server.




Figure 2. Location on the BladeCenter server planar where the CFFh card is installed

IBM BladeCenter chassis support is based on the blade server type in which the expansion card is installed. Consult ServerProven to see which chassis each blade server type is supported in: <http://ibm.com/servers/eserver/serverproven/compat/us/>.

Table 3 lists the I/O modules that can be used to connect to the 2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter. The I/O modules listed in Table 3 are supported in BladeCenter H chassis only.

Table 3. I/O modules supported with the 2-Port 40 Gb InfiniBand Expansion Card (CFFh)

I/O module	Part number	BladeCenter S	BladeCenter E	BladeCenter H	BladeCenter T	BladeCenter HT	MSIM	MSIM-HT
Voltaire 40 Gb InfiniBand Switch Module	46M6005	N	N	Y	N	N	N	N

In BladeCenter H, the ports of CFFh cards are routed through the midplane to I/O bays 7, 8, 9, and 10, as shown in Figure 3.

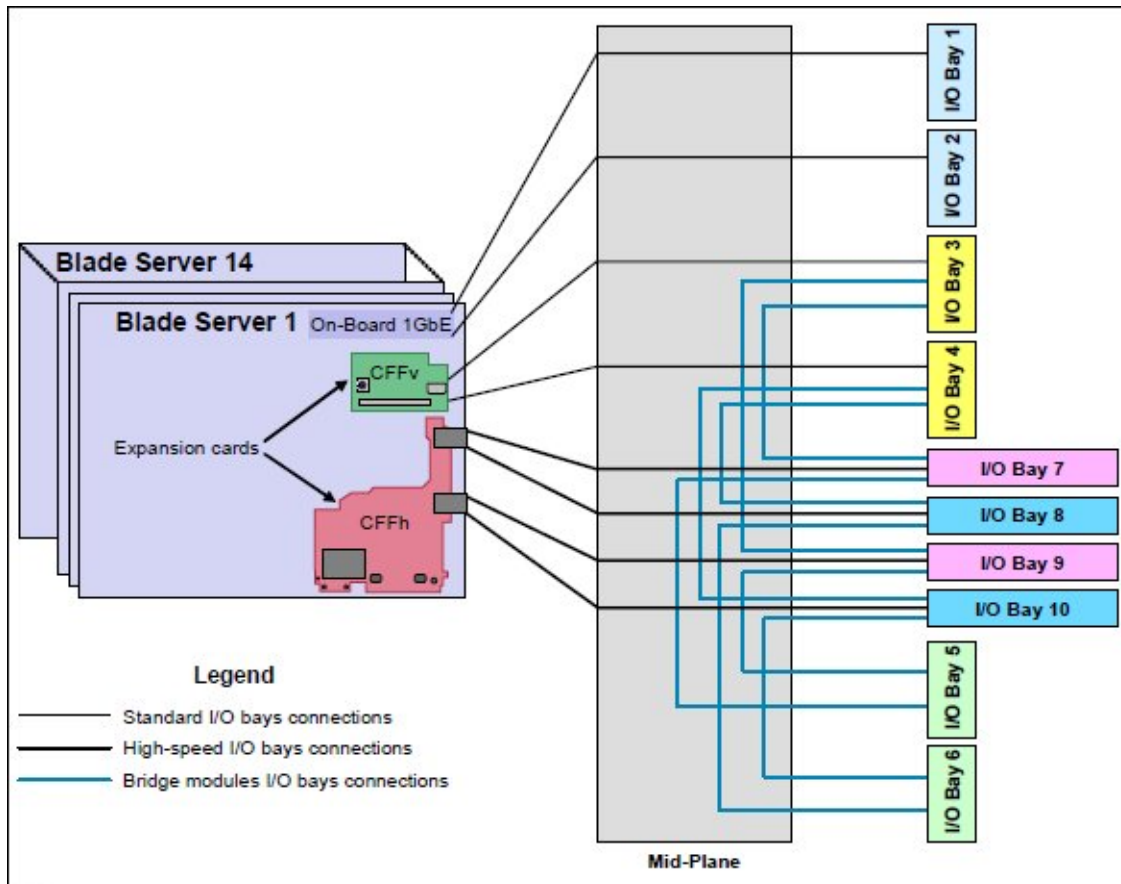


Figure 3. IBM BladeCenter H I/O topology showing the I/O paths from CFFh expansion cards

One I/O module must be installed in the chassis for each 4X InfiniBand port that you wish to use on the expansion card. The specific I/O bays in the chassis are listed in Table 4. For the 2-Port 40 Gb InfiniBand Expansion Card (CFFh), you should install an I/O module in I/O bays 7/8 and 9/10 (that is, two I/O modules, each of them occupies two adjacent high-speed slots).

Table 4. Locations of I/O modules required to connect to the expansion card

Expansion card	I/O bay 7	I/O bay 8	I/O bay 9	I/O bay 10
2-Port 40 Gb InfiniBand Expansion Card (CFFh)	Supported I/O module*		Supported I/O module*	

* A single Voltaire 40 Gb InfiniBand Switch Module occupies two adjacent high-speed bays (7 and 8 or 9 and 10) while expansion cards have only two ports--one port per one InfiniBand module.

Popular configurations

Figure 4 shows the use of Voltaire 40 Gb InfiniBand Switch Module to route two 4X InfiniBand ports from 2-Port 40 Gb InfiniBand Expansion Card (CFFh) installed into each server. Two Voltaire 40 Gb InfiniBand Switch Modules are installed in bays 7/8 and bays 9/10 of the BladeCenter H chassis. All connections between the expansion cards and the switch modules are internal to the chassis. No cabling is needed.




Figure 4. A 40 Gb solution using 2-Port 40 Gb InfiniBand Expansion Card (CFFh) and Voltaire 40 Gb InfiniBand Switch Modules

Table 5 lists the components that this configuration uses.

Table 5. Components used when connecting 2-Port 40 Gb InfiniBand Expansion Card (CFFh) to two Voltaire 40 Gb InfiniBand Switch Modules

Diagram reference	Part number/machine type	Description	Quantity
1	Varies	IBM BladeCenter HS22 or other supported server	1 to 14
2	46M6001	2-Port 40 Gb InfiniBand Expansion Card (CFFh)	1 per server
3	8852	BladeCenter H	1
4	46M6005	Voltaire 40 Gb InfiniBand Switch Module	2
5	49Y9980	3 m Copper QDR InfiniBand QSFP Cable	Up to 32*

* The Voltaire 40 Gb InfiniBand Switch Module has 16 external ports. To communicate outside of the chassis, you must have QSFP cables connected. You have the flexibility to expand bandwidth using from one to 16 connections per switch.

Operating system support

The 2-Port 40 Gb InfiniBand Expansion Card (CFFh) supports the following operating systems:

- Microsoft Windows Server 2008 R2
- Microsoft Windows Server 2008, Datacenter x64 Edition
- Microsoft Windows Server 2008, Datacenter x86 Edition
- Microsoft Windows Server 2008, Enterprise x64 Edition
- Microsoft Windows Server 2008, Enterprise x86 Edition
- Microsoft Windows Server 2008, Standard x64 Edition
- Microsoft Windows Server 2008, Standard x86 Edition
- Microsoft Windows Server 2008, Web x64 Edition
- Microsoft Windows Server 2008, Web x86 Edition
- Red Hat Enterprise Linux 4 AS for AMD64/EM64T
- Red Hat Enterprise Linux 4 AS for x86
- Red Hat Enterprise Linux 5 Server Edition
- SUSE LINUX Enterprise Server 10 for AMD64/EM64T
- SUSE LINUX Enterprise Server 11 for AMD64/EM64T

Support for operating systems is based on the combination of the expansion card and the blade server in which it is installed. See IBM ServerProven for the latest information about the specific versions and service packs supported: <http://ibm.com/servers/eserver/serverproven/compat/us/>. Select the blade server, and then select the expansion card to see the supported operating systems.

Related publications

For more information, refer to the following documents:

- *2-Port 40 Gb InfiniBand Expansion Card (CFFh) for IBM BladeCenter Installation and User's Guide*
<http://www.ibm.com/support>
- *Voltaire 40 Gb InfiniBand Switch Module Installation Guide*
<http://www.ibm.com/support>
- IBM BladeCenter Interoperability Guide
<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073016>
- IBM Redbooks publication *IBM BladeCenter Products and Technology*, SG24-7523
<http://www.redbooks.ibm.com/abstracts/sg247523.html>

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.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2009. All rights reserved.

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

This document was created or updated on December 7, 2010.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips0700.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>. The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®
IBM®
Redbooks®
Redpaper™
Redbooks (logo)®
ServerProven®
System x®

The following terms are trademarks of other companies:

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

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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