

Voltaire 40 Gb InfiniBand Switch Module for IBM BladeCenter

IBM BladeCenter at-a-glance guide

Clients want to increase the use of their existing servers to reduce overall cost as well as the IT footprint within the organization. IBM BladeCenter offers several choices, including servers, chassis, and the types of interconnect fabric. This solution offers end-to-end 40 Gb (4X QDR) InfiniBand connectivity. Clients can take advantage of the resiliency of BladeCenter as well as the performance of 4X QDR InfiniBand in a single package.

The Voltaire 40 Gb InfiniBand Switch Module for BladeCenter provides InfiniBand QDR connectivity between the server platform's high performance blade servers and external InfiniBand fabrics in non-blocking designs, all on a single device. Voltaire's high speed module also accommodates performance-optimized fabric designs utilizing a single BladeCenter chassis or stacking multiple BladeCenter chassis without requiring an external InfiniBand switch.

The InfiniBand switch module offers 14 Gb ports, one to each server, and 16 ports out of the chassis per switch. This solution offers a no compromise, congestion free solution to meet even the most performance hungry applications.

Figure 1 shows the switch module.



Figure 1. Voltaire 40 Gb InfiniBand Switch Module for IBM BladeCenter

Did you know?

The module's HyperScale architecture also provides a unique inter-switch link or mesh capability to form highly scalable, cost effective, and low latency fabrics. Since this switch has 16 uplink ports, they can create a meshed architecture and still have unblocked access to data using the 14 uplink ports. This solution can scale from 14 to 126 nodes and offer latency less than 200 nanosecond allowing applications to operate at maximum efficiency.

Part number information

Table 1 shows the part numbers to order the module and the supported cables.

Table 1. Part number and feature code for ordering

Description	Part number	Feature codes
Voltaire 40 Gb InfiniBand Switch Module for IBM BladeCenter	46M6005	0057
3 m Copper QDR InfiniBand QSFP Cable	49Y9980	3866

The Voltaire 40 Gb InfiniBand Switch Module part number includes the following items:

- One Voltaire 40 Gb InfiniBand Switch Module
- Documentation CD-ROM
- Important Notices flyer

Features and specifications

The Voltaire 40 Gb InfiniBand Switch Module includes the following features and functions:

- Form-factor
 - Double-height high-speed switch module
- Internal ports
 - 14 internal 4X QDR/DDR/SDR InfiniBand to the server blades
 - Two internal 4X InfiniBand ports to the bridge modules
- External ports
 - Up to 16 auto-sensing 4X QDR/DDR/SDR InfiniBand QSFP ports (40, 20, or 10 Gbps auto-negotiate)
- Characteristics
 - Full QDR rate InfiniBand switching
 - Based on the Infiniscale-IV device
 - Up to 40 Gbps performance for clusters and grids (bidirectional)
 - Hot swap support, with soft start and current limiting
 - Bay address and presence support
 - High temperature monitoring
 - Full BladeCenter H AMM interface via I2C
 - Fully non-blocking QDR blade
 - QSFP front panel connectors
 - On board CPU for chassis management capabilities
 - Ultra-low latency: under 100 nanoseconds blade to blade connectivity
 - Built-in high availability
 - Port status controllability

Supported BladeCenter chassis and expansion cards

The Voltaire 40 Gb InfiniBand Switch Module is supported only in the IBM BladeCenter H chassis as listed in Table 2.

Table 2. IBM BladeCenter chassis that support the Voltaire 40 Gb InfiniBand Switch Module

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

The Voltaire 40 Gb InfiniBand Switch Module support the expansion cards listed in Table 3. This table also lists the chassis bays in which the switch module must be installed when used with each expansion card.

Table 3. Voltaire 40 Gb InfiniBand Switch Module and BladeCenter chassis I/O bays support

Description	Part number	Bay 1 (Standard)	Bay 2 (Standard)	Bay 3 (Standard)	Bay 4 (Standard)	Bay 5 (Bridge)	Bay 6 (Bridge)	Bay 7 (High-speed)	Bay 8 (High-speed)	Bay 9 (High-speed)	Bay 10 (High-speed)
2-Port 40Gb InfiniBand Expansion Card (CFFh)	46M6001	N	N	N	N	N	N	Y	Y	Y	Y
InfiniBand 4X DDR Expansion Card (CFFh)	43W4423	N	N	N	N	N	N	Y	Y	Y	Y
Voltaire 4X InfiniBand DDR (CFFh)	43W4420	N	N	N	N	N	N	Y	Y	Y	Y

Note: The Voltaire 40 Gb InfiniBand Switch Module is a double-high switch module and occupies two adjacent high-speed bays (7 and 8 or 9 and 10).

The BladeCenter H chassis has four standard I/O bays (1, 2, 3, and 4), two bridge bays (5 and 6), and four high-speed bays (7, 8, 9, and 10). The Voltaire 40 Gb InfiniBand Switch Module is supported only in the BladeCenter H chassis and occupies two adjacent high-speed I/O bay (bays 7 and 8, or bays 9 and 10).

Popular configurations

Figure 2 shows the use of Voltaire 40 Gb InfiniBand Switch Module to route two 4X InfiniBand ports from 2-Port 40Gb 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 2. A 40 Gb solution using two Voltaire 40 Gb InfiniBand Switch Modules

Table 4 lists the components that this configuration uses.

Table 4. 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.

Connectors and LEDs

Figure 3 shows the left side of the Voltaire 40 Gb InfiniBand Switch Module front panel.




Figure 3. Left side of front panel of the Voltaire 40 Gb InfiniBand Switch Module*

* The port LEDs on the right side of front panel of the Voltaire 40 Gb InfiniBand Switch Module are the same as on the left side.

The front panel contains the following components:

- Three status LEDs display the status of the switch module and the network:
 - OK (indicating that the switch module has passed the power-on self-test (POST) with no critical faults and is operational)
 - Switch-module error (indicating that the switch module has failed the POST or detected an operational fault)
 - SM (reserved for future use)
- 16 QSFP port connectors to attach QSFP cables.
 - Each external QSFP port on the switch module contains IB physical link and IB logical link LEDs

Network cabling requirements

The Voltaire 40 Gb InfiniBand Switch Module supports the following cable:

- 3 m Copper QDR InfiniBand QSFP Cable, part number 49Y9980

Related publications

For more information, see the following documents:

- IBM US Announcement Letter for the Voltaire 40 Gb InfiniBand Switch Module
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS109-377>
- *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 4, 2009.

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/tips0696.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 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®
Redbooks (logo)®
System x®

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