

Overview

HP supports 40Gbps (QDR) and 20Gbps (DDR) InfiniBand products that include mezzanine Host Channel Adapters (HCA) for server blades, switch blades for c-Class enclosures, and rack switches and cables for building scale-out solutions. This QuickSpecs focuses on mezzanine HCAs for server blades, and InfiniBand switch blades for c-Class enclosures; for details on the InfiniBand rack switches, standup PCI Express HCAs for HP ProLiant and Integrity servers, as well as InfiniBand cables, please refer to the HP InfiniBand for HP ProLiant and Integrity servers QuickSpecs at:

http://h18000.www1.hp.com/products/quickspecs/13078_div/13078_div.html

HP supports the following InfiniBand products for HP BladeSystem c-Class

- HP 4X QDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class
- HP 4X DDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class
- HP 4X DDR IB (single-port) Mezzanine HCA for HP BladeSystem c-Class
- HP BLc 4X QDR IB Switch for HP BladeSystem c-Class
- HP BLc 4X DDR IB G2 Switch for HP BladeSystem c-Class
- HP 4X DDR IB Switch Module for HP BladeSystem c-Class

The HP 4X QDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class is based on the Mellanox ConnectX IB technology. The QDR IB HCA delivers low-latency and up to 40Gbps (QDR) bandwidth for performance-driven server and storage clustering applications in High-Performance Computing (HPC) and enterprise data centers. Parallel or distributed applications running on multi-processor multi-core servers will benefit from the reliable transport connections and advanced multicast support offered by ConnectX IB. End-to-end Quality of Service (QoS) enables partitioning and guaranteed service levels while hardware-based congestion control prevents hot spots from degrading the effective throughput. The HP 4X QDR IB Dual-port Mezzanine HCA card is designed for PCI Express 2.0 x8 connectors on HP BladeSystem c-Class G6 server blades.

The HP 4X DDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class is based on the Mellanox ConnectX IB technology. The DDR IB HCA delivers low-latency and up to 20Gbps (DDR) bandwidth for performance-driven server and storage clustering applications in High-Performance Computing (HPC) and enterprise data centers. The HP 4X DDR Dual-port Mezzanine HCA card is designed for PCI Express connectors on HP BladeSystem c-Class server blades. To obtain the best performance, it is recommended to plug the card into the x8 PCI Express connector on the server blades.

The HP 4X DDR IB (single-port) Mezzanine HCA for HP BladeSystem c-Class is a single-port card, capable of supporting up to 20Gbps (DDR) bandwidth. The HP 4X DDR IB single-port HCA card is designed for PCI Express connectors on HP BladeSystem c-Class server blades. To obtain the best performance, it is recommended to plug the card into the x8 PCI Express connector on the server blades.

InfiniBand host stack software (driver) is required to run on servers connected to the InfiniBand network. HP supports Voltaire OFED driver stack on Linux 64-bit operating systems, and WinOF driver stack on Microsoft Windows (HPC) server 2008. A right-to-use (RTU) license is required for Voltaire OFED driver stack.

Voltaire OFED includes the upper level protocol (ULP) for:

- IPoIB-UD and IPoIB-CM
- SDP
- Bonding driver
- iSER
- uDAPL
- RDS
- NFSoRDMA - technology preview - supported only on kernel 2.6.27

WinOF includes the ULP for:

Overview

- IPoB
- Winsock Direct (WSD)
- NetworkDirect (ND)

HP-MPI is supported on Voltaire OFED, and WinOF. HP-MPI is a high performance and production quality implementation of the Message-Passing Interface (MPI) standard for HP servers, blades and workstations. HP-MPI uses enhancements whenever appropriate to provide low latency and high bandwidth point-to-point and collective communication routines. It supports multi-protocol execution of MPI applications on clusters of shared-memory servers so that applications can take advantage of the shared memory for intra-node communications.

The HP BLC 4X QDR IB Switch for HP BladeSystem c-Class is a double wide switch for the new HP BladeSystem c7000 enclosure. It is based on the Mellanox InfiniScale IV technology. The QDR IB switch blade has 16 downlink ports to connect up to 16 server blades in the c7000 enclosure, and 16 QSFP uplink ports for inter-switch links or to connect to external servers. All ports are capable of supporting 40Gbps (QDR) bandwidth. A subnet manager has to be provided; see the paragraph on subnet managers for more details. The HP BLC 4X DDR IB Gen2 Switch is a 2nd generation InfiniBand double wide switch for the HP BladeSystem c7000 and c3000 enclosures. It is based on the Mellanox InfiniScale IV technology. The DDR Gen2 IB switch has 16 downlink ports to connect up to 16 server blades in the enclosures, and 16 QSFP uplink ports for inter-switch links or to connect to external servers. All ports are capable of supporting 20Gbps (DDR) bandwidth. A subnet manager has to be provided; see the paragraph on subnet managers for more details.

The HP 4X DDR IB Switch Module for HP BladeSystem c-Class is a double wide switch for the HP BladeSystem c7000 and c3000 enclosures. It is based on the Mellanox InfiniScale III technology. The HP 4X DDR IB Switch Module blade has 16 downlink ports to connect up to 16 server blades in the enclosures, and 8 CX4 uplink ports for inter-switch links or to connect to external servers. All ports are capable of supporting 20Gbps (DDR) bandwidth. A subnet manager has to be provided; see the paragraph on subnet managers for more details.

An InfiniBand fabric consists of one or more InfiniBand switches connected via inter-switch links. The most commonly deployed fabric topology is a fat tree or its variations. A subnet manager is required to manage and control an InfiniBand fabric. The subnet manager functionality can be provided by either a rack-mount InfiniBand switch with an embedded fabric manager (aka internally managed switch) or host-based subnet manager software on a server connected to the fabric. Embedded fabric manager is available on Voltaire internally managed 24-port, 96-port, and 288-port DDR switches, and 36-port QDR switch. Please refer to http://h18000.www1.hp.com/products/quickspecs/13078_div/13078_div.html for information about Voltaire InfiniBand rack switches. Voltaire HBSM is a host-based subnet manager that supports Voltaire OFED on Linux operating system and supersedes Voltaire GridVision BladeFM. OpenSM for Windows is a host-based subnet manager that supports WinOF on Microsoft Windows HPC server 2008.

HP supports InfiniBand copper and fiber optic cables with CX4 to CX4, CX4 to QSFP, and QSFP to QSFP connectors. The CX4 to CX4 copper cables range from 0.5M to 8M for HCA to switch, or inter-switch links at DDR speed, and up to 12M for certain inter-switch links at DDR speed. The CX4 to CX4 fiber optic cables range from 1M to 100M for HCA to switch, or inter-switch links at DDR speed. Please note that not every InfiniBand switch port supports fiber optic cables. The CX4 to QSFP copper cables range from 1M to 5M for HCA to switch, or inter-switch links at either DDR or QDR speed. The QSFP to QSFP copper cables range from 1M to 5M for HCA to switch, or inter-switch links at either DDR or QDR speed, and up to 10M at DDR speed. Please refer to http://h18000.www1.hp.com/products/quickspecs/13078_div/13078_div.html for more details on supported InfiniBand cables.

What's New

- HP 4X QDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class G6 server blades
 - HP BLC 4X QDR InfiniBand (IB) Switch for new HP BladeSystem c7000 enclosure
-

Overview

At A Glance

- HP 4X QDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class
 - Dual-port 4X QDR InfiniBand PCI Express G2 Mezzanine card
 - Supported on HP ProLiant BL460c G6 and BL490c G6 c-Class server blades
 - Support Voltaire OFED Linux driver stack and WinOF 2.0 on Microsoft Windows HPC server 2008
- The HP 4X DDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class
 - Dual-port 4X DDR InfiniBand PCI Express Mezzanine card
 - Supported on HP ProLiant BL260c G5, BL2x220c G5 Server, BL460c, BL460c G5, BL460c G6, BL465c G5, BL480c, BL490c G6, BL495c G5, BL680c G5, BL685c, BL685c G5, BL685c G6 c-Class server blades, and HP Integrity BL860c
 - Support Voltaire OFED Linux driver stack and WinOF 2.0 on Microsoft Windows HPC server 2008 (HP ProLiant blades only)
- The HP 4X DDR IB Mezzanine HCA for HP BladeSystem c-Class
 - Single-port 4X DDR InfiniBand PCI Express Mezzanine card
 - Supported on HP ProLiant BL260c G5, BL2x220c G5, BL460c, BL460c G5, BL460c G6, BL465c G5, BL480c, BL490c G6, BL495c G5, BL680c G5, BL685c, BL685c G5, BL685c G6 c-Class server blades, and HP Integrity BL860c, with dual-core processors
 - Support Voltaire OFED Linux driver stack and WinOF 2.0 on Microsoft Windows HPC server 2008 (HP ProLiant blades only)
- HP BLc 4X QDR IB Switch for HP BladeSystem c-Class
 - Double wide switch blade for the new c7000 enclosure with 16 downlink ports to connect server blades via the midplane, and 16 QSFP uplink ports for the inter-switch links or to connect to external servers. All ports are capable of supporting 40Gbps (QDR) bandwidth. All uplink ports support copper and fiber optic cables. The switch is externally managed, i.e.,
 - subnet manager has to be provided on the fabric (see the subnet manager discussion above)
- The HP BLc 4X DDR IB Gen2 Switch for HP BladeSystem c-Class
 - Double wide switch module for c-Class enclosures, with 16 downlink ports to connect server blades via the midplane, and 16 QSFP uplink ports for the inter-switch links or to connect to external servers. All ports are capable of supporting 20Gbps DDR bandwidth. All uplink ports support copper and fiber optic cables. The switch is externally managed, i.e., a subnet manager has to be provided on the fabric (see the subnet manager discussion above)
- The HP 4X DDR IB Switch Module for HP BladeSystem c-Class
 - Double wide switch module for c-Class enclosures, with 16 downlink ports to connect server blades via the midplane, and 8 CX4 uplink ports for the inter-switch links or to connect to external servers. All ports are capable of supporting 20Gbps DDR bandwidth. All uplink ports support copper and fiber optic cables. The switch is externally managed, i.e., a subnet manager has to be provided on the fabric (see the subnet manager discussion above)

Models

HP IB Mezz HCA cards for BladeSystem c-Class	HP 4X QDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class	492303-B21
	HP 4X DDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class	448262-B21
	HP 4X DDR IB Mezzanine HCA for HP BladeSystem c-Class (single-port)	410533-B21
	Voltaire OFED driver stack RTU license	450716-B21
HP InfiniBand Switch for BladeSystem c-Class and Subnet Manager Option	HP BLc 4X QDR IB Switch for HP BladeSystem c-Class	489184-B21
	HP BLc 4X DDR IB Gen2 Switch for HP BladeSystem c-Class	489183-B21
	HP 4X DDR IB Switch Module for HP BladeSystem c-Class	410398-B21
	Voltaire Host-based Subnet Manager (HBSM) for up to two enclosures (up to 64 HCAs). Voltaire HBSM supersedes Voltaire GridVision BladeFM	450717-B21

Performance

Latency and Bandwidth

All ports on 4X QDR HCA card and switch are capable of supporting 40Gbps signaling rate, with a peak data rate of 32 Gbps in each direction.

All ports on 4X DDR HCA cards and switches are capable of supporting 20 Gbps signaling rate, with a peak data rate of 16 Gbps in each direction.

Scalability and Reliability

Standards support

- Dual-port 4X QDR mezzanine HCA: PCI Express revision 2.0x8 (1.1 compliant)
- Dual-port 4X DDR mezzanine HCA: PCI Express revision 2.0x8 (1.1 compliant)
- Single-port 4X DDR Mezzanine HCA: PCI Express revision 1.0a
- IBTA version 1.2 compatible

Operating systems support

- Voltaire OFED for Linux:
 - RHEL 4 U4, 5, 6, 7, and RHEL 5 U1, 2
 - SLES 10 SP1, 2, 3
- WinOF for Microsoft Windows HPC server 2008 (64-bit)
 - Microsoft Windows HPC server 2008

PCI-Express Mezzanine connectors

The HP 4X QDR IB HCA mezzanine card is specified to run on PCI Express Gen2 connectors.

The HP 4X DDR IB HCA mezzanine cards are specified to run on PCI Express mezzanine connectors. Either x8 or x4 connectors can be used for the IB 4X DDR HCA cards, however, for the best 4X DDR performance, DDR IB HCA mezzanine cards should only be plugged on x8 PCI Express connectors.

Depending on the server blade models, multiple HCAs per server blade may be supported; please refer to server blades specification for more details.

Switch bays on c-Class enclosure

HP IB switch blades are designed to fit into the double wide switch bays on the c-Class enclosures. Depending on the mezzanine connectors used for IB HCA mezzanine cards, the IB switch blade has to be inserted into corresponding switch bays.

InfiniBand Fabric Management

Auto-negotiation/ self-discovery

InfiniBand QDR ports support auto-negotiation down to DDR/SDR, InfiniBand DDR ports support auto-negotiation down to SDR. For example, external uplink ports of a DDR switch blade will negotiate down to SDR speed if they are connected to a device which can only operate at SDR.

Self-discovery is supported by subnet manager. Check the documents of the specific subnet manager for more details (<http://www.docs.hp.com>).

Management Support

Voltaire® OFED includes drivers and utilities to configure and manage the HCA in a Linux environment and WinOF includes drivers and utilities to configure and manage the HCA in the Microsoft Windows server HPC 2008 environment.

InfiniBand fabric has to be set up with a subnet manager: see the subnet manager discussion above.

Additional Features

HP Services and Support The HP 4X QDR/DDR IB HCA cards include one year parts-only exchange or covered by the warranty of the server installed in whichever is greater.

The HP 4X QDR/DDR IB switches include one year parts only exchange warranty. Enhancements to warranty services are available through HP Care Pack services or customized service agreements. The following Flexible CarePack services are available for the IB switch products:

- 1, 3, 4, 5 year Next Business Day HW Support
- 1, 3, 4, 5 year Same Day HW Support
- 1, 3, 4, 5 year 24x7 HW Support
- 1, 3, 4, 5 year 6 hour Call to Repair HW Support
- 1, 3, 4, 5 year Support Plus Service
- 1, 3, 4, 5 year Support Plus 24 Service
- 1, 3, 4, 5 year Proactive 24 Service
- 1, 3, 4, 5 year Critical Service

NOTE: For more information, visit HP Care Pack Services at: <http://www.hp.com/services>.

Related Options

The HP BLc 4X QDR IB Switch is only supported on the c7000 enclosure with part numbers: 507014-B21, 507015-B21, 507016-B21, 507017-B21, 507019-B21.

Technical Specifications

Compliance	Mezzanine:	Dual-port QDR HCA: PCI Express 2.0 Dual-port DDR HCA: PCI Express 2.0 (1.1 compliant) Single-port DDR HCA: PCI Express revision 1.0a IBTA version 1.2 compatible ROHS-R5
	Switches:	IBTA version 1.2 compatible ROHS-R5
General Specifications	Communications Processor	Mezzanine: Dual-port QDR HCA: MT25408-FCC-QI Dual-port DDR HCA: MT25408A0-FCC-D Single-port DDR HCA : MT25204A0-FCC-D
		QDR switch: MT48436A1-FCC-Q, InfiniScale IV DDR Gen2 switch: MT48436A1-FCC-D, InfiniScale IV DDR switch module: MT47396A1-FDC-D, InfiniScale III
	Dimensions (LxW)	Mezzanine HCA: 4.5 x 4.0 in (11.43 x 10.16 cm) Switches: 15.3 x 10.6 in (38.86 x 26.92 cm)
	Power and Environmental Specifications	Operating Temperature Mezzanine HCA: 32° to 131° F (0° to 55° C) Switch module: 32° to 104° F (0° to 40° C)
	Humidity (non-condensing) Mezzanine HCA: 5% to 85% Switch module: 5% to 85%	
	Power requirement Mezzanine HCAs: Dual-port QDR HCA: max 13.5W Dual-port DDR HCA: max 12.9W Single-port DDR HCA: max 4.5W Switch blades: QDR switch: 10.8A at 12V Max (130W) DDR Gen2 switch: 10 A at 12V max (120 W) DDR switch module: 4.0 A at 12V max (48 W)	
	Emissions Classifications HCA and Switches: FCC CFR 47 Part 15 Class A CISPR 22 Class A ICES-003 Class A VCCI Class A ACA CISPR 22 Class A Switches: MIC Korea "EMC Registration Regulation" Class A	

© Copyright 2009 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice.

Windows and Microsoft are registered trademarks of Microsoft Corp., in the U.S.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein