

Overview

HP supports 40Gbps 4x Quad Data Rate (QDR) and 20Gbps 4X Double Data Rate (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_na/13078_na.html

The following InfiniBand products based on Mellanox technology are available for the HP BladeSystem c-Class from HP:

- HP 4X QDR IB CX-2 Dual Port Mezz HCA for HP BladeSystem c-Class
- HP 4X DDR IB Dual-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

The HP 4X QDR IB CX-2 Dual-port Mezzanine HCA for HP BladeSystem c-Class is based on the Mellanox ConnectX-2 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 CX-2 Dual-port Mezzanine HCA card is designed for PCI Express 2.0 x8 connectors on HP BladeSystem c-Class G6 server blades. Depending on the server blade models, some configuration rules may apply which may result in only one port being connected to the enclosure mid-plane. Please refer to server blades specification for more details.

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. Depending on the server blade models, some configuration rules may apply which may result in only one port being connected to the enclosure mid-plane. Please refer to server blades specification for more details.

InfiniBand host stack software (driver) is required to run on servers connected to the InfiniBand network. For mezzanine HCAs based on Mellanox technologies, HP supports Mellanox OFED and Voltaire OFED driver stacks on Linux 64-bit operating systems, and Mellanox WinOF driver stack on Microsoft Windows (HPC) server 2008.

A right-to-use (RTU) license is required for Voltaire OFED driver stack. For the latest information on Voltaire OFED software, please refer to: http://www.voltaire.com/Products/Application_Acceleration_Software/Voltaire_OFED_and_WinOF_Support.

For the latest information on Mellanox OFED software, please refer to: http://www.mellanox.com/content/pages.php?pg=software_overview_ib&menu_section=34.

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 double wide switch for the HP BladeSystem c7000 and c3000 enclosures. It is based on the

Overview

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.

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. OpenSM is a host-based subnet manager that runs on a server connected to the InfiniBand fabric. Mellanox OFED software stack includes OpenSM for Linux, and Mellanox WinOF includes OpenSM for Windows. For comprehensive management and monitoring capabilities, Mellanox FabricIT™ is recommended for managing the InfiniBand fabric based on Mellanox InfiniBand products; and Voltaire Unified Fabric Manager™ (UFM) is recommended for managing the InfiniBand fabric based on Voltaire InfiniBand switch products and Mellanox ConnectX-based mezzanine HCA running Voltaire OFED stack.

Embedded fabric manager is available on Voltaire internally managed 36-port 4X QDR switch, 24-port, 96-port, and 288-port DDR switches. Please refer to: http://h18000.www1.hp.com/products/quickspecs/13078_na/13078_na.html for information about InfiniBand rack switches.

The following InfiniBand products based on QLogic technology are available for the HP BladeSystem c-Class from HP:

- QLogic 4X QDR IB Mezzanine HCA for HP BladeSystem c-Class
- QLogic BLc 4X QDR IB Switch for HP BladeSystem c-Class
- QLogic BLc 4X QDR IB Management Module

The QLogic 4X QDR IB Mezzanine HCA is based on the QLogic TrueScale ASIC architecture which has a unique hardware architecture that delivers unprecedented levels of performance, reliability, and scalability, making it an ideal solution for highly scaled High Performance Computing (HPC) and high throughput, low-latency enterprise applications.

InfiniBand host stack software (driver) is required to run on servers connected to the InfiniBand fabric. For HCAs based on QLogic technology, HP supports QLogic OFED driver stacks on Linux 64-bit operating systems.

The QLogic BLc 4X QDR IB Switch for HP BladeSystem c-Class uses the QLogic TrueScale ASIC architecture designed to cost-effectively link workgroup resources into a cluster or provide an edge switch option for a larger fabric. Customers have the ability to internally or externally manage the modular IB switch. The QLogic BLc 4X QDR IB Switch also has an optional management module that includes an embedded subnet manager. This management module is an option for both the 32-port QDR BladeSystem switch and the 36-port QDR edge switch. When combined with the optional InfiniBand Fabric Suite software, users can manage up to a 288 node fabric using only the management capability of the unit, without requiring additional host processing. In a bladed environment this alleviates the need to use a node for fabric management.

An InfiniBand fabric is constructed with 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 an InfiniBand fabric. OpenSM is a host-based subnet manager that runs on a server connected to the InfiniBand fabric. QLogic OFED software stack includes OpenSM for Linux. For comprehensive management and monitoring capability, QLogic InfiniBand Fabric Suite (IFS) is recommended for managing the InfiniBand fabric based on QLogic InfiniBand products.

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 only 12 ports on Voltaire 24-port DDR switches support 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 7M for HCA to switch, or inter-switch links at either DDR or QDR speed (please note that QLogic QDR switches only support up to 5 meters at QDR speed) , and up to 10M at DDR speed. Please refer to:

Overview

http://h18000.www1.hp.com/products/quickspecs/13078_na/13078_na.html for more details on supported InfiniBand cables.

What's New

- HP 4X QDR IB CX-2 Dual Port Mezz HCA for HP BladeSystem c-Class
- QLogic 4X QDR IB Dual Port Mezzanine HCA for HP BladeSystem c-Class
- QLogic BLc 4X QDR IB Switch for HP BladeSystem c-Class

At A Glance

- **InfiniBand mezzanine HCA based on Mellanox technology**
 - HP 4X QDR IB CX-2 Dual-port Mezzanine HCA for HP BladeSystem c-Class
 - Dual-port 4X QDR InfiniBand PCI Express G2 Mezzanine card
 - Supported on HP ProLiant BL280c G6, BL2x220c G6, BL460c G6 and BL490c G6 c-Class server blades
 - Support Mellanox OFED and Voltaire OFED Linux driver stacks and Mellanox WinOF 2.0 on Microsoft Windows HPC server 2008
 - 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, BL280c G6, BL2x220c G5 and G6 , 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 Mellanox OFED and Voltaire OFED Linux driver stack and Mellanox WinOF 2.0 on Microsoft Windows HPC server 2008 (HP ProLiant blades only)
- **InfiniBand mezzanine HCA based QLogic technology**
 - QLogic 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 BL280c G6, BL2x220c G6, BL460c G6, BL490c G6 c-Class server blades
 - Full support for QLogic OFED including all standard OFED transport protocols, plus optional value-added components to improve management and performance
- **InfiniBand switch blades based on Mellanox technology**
 - 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 unlink 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. a subnet manager has to be provided on the fabric (see the subnet manager discussion above)
 - 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)
- **InfiniBand switch blades based on QLogic technology**
 - QLogic 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 unlink 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
 - Optional Management Module with an embedded subnet manager and management of up to 288 nodes

Models

InfiniBand mezzanine HCA cards and options	NOTE: The HCA cards listed in this section are based on Mellanox technology.	
	HP 4X QDR IB CX-2 Dual Port Mezz HCA for HP BladeSystem c-Class	592519-B21
	HP 4X DDR IB Dual-port Mezzanine HCA for HP BladeSystem c-Class	448262-B21
	Voltaire OFED driver stack RTU license	450716-B21
	Voltaire Unified Fabric Mgr (UFM) Per Node RTU	590643-B21
	NOTE: Right to use (RTU) License for Voltaire Unified Fabric Manager (UFM), 1 per node for every node connected to the Voltaire InfiniBand fabric.	
QLogic IB mezzanine HCA card	QLogic 4X QDR IB Dual Port Mezz HCA	583210-B21
Mellanox IB switch blades and options	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
	NOTE: The HP BLc 4X QDR IB Switch is only supported on the HP BladeSystem c7000 Enclosures (P/N's: 507014-B21, 507015-B21, 507016-B21, 507017-B21 and 507019-B21).	
	NOTE: Please see the HP BladeSystem c7000 Enclosure QuickSpecs for additional information at: http://h18000.www1.hp.com/products/quickspecs/12810_na/12810_na.html	
	Mellanox Fabric IT Mgr Per Node RTU	588055-B21
	NOTE: Right to use (RTU) License for Mellanox FabricIT, 1 per node for every node connected to the Mellanox InfiniBand fabric.	
	Voltaire Unified Fabric Mgr Per Node RTU	590643-B21
	NOTE: Right to use (RTU) License for Voltaire Unified Fabric Manager (UFM), 1 per node for every node connected to the Voltaire InfiniBand fabric.	
QLogic IB switch blade and options	QLogic BLc 4X QDR IB Switch for HP BladeSystem c-Class	505958-B21
	QLogic BLc 4X QDR IB Management Module	505959-B21
	QLogic InfiniBand Fabric Suite (IFS)	589485-B21
	NOTE: Right to use (RTU) License for QLogic IFS, 1 per node for every node connected to the QLogic InfiniBand fabric.	

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)
- IBTA version 1.2 compatible

Operating systems support

- Mellanox OFED for Linux
 - RHEL 5, RHEL 5 U1, RHEL 5 U2, SLES10, SLES10 SP1, SLES10 SP1 (up1), SLES 10 SP2, RHEL 4 U4, U5, U6,
- Mellanox WinOF for Microsoft Windows HPC server 2008 (64-bit)
 - Microsoft Windows HPC server 2008
- Voltaire OFED for Linux:
 - RHEL 4 U4, 5, 6, 7, and RHEL 5 U1, 2
 - SLES 10 SP1, 2, 3
- QLogic OFED+ for Linux (64-bit)
 - RHEL 5.2, 5.3
 - SLES 10 SP2, 11
 - CentOS (Rocks) 5.2, 5.3
 - Scientific Linux 5.2, 5.3

PCI-Express Mezzanine connectors

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

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

Mellanox OFED, Voltaire® OFED and QLogic OFED+ include drivers and utilities to configure and manage the HCA in a Linux environment and Mellanox 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

Warranty

HP branded hardware options qualified for BladeSystem c-Class and p-Class servers are covered by a global limited warranty and supported by HP Services and a worldwide network of HP Authorized Channel Partners. The HP branded hardware option diagnostic support and repair is available for one year from date of purchase, or the length of the server they are attached to, whichever is greater. Support for software and initial setup is available for 90 days from date of purchase. Additional support may be covered under the warranty or available for an additional fee. Enhancements to warranty services are available through HP Care Pack services or customized service agreements.

Additional information regarding worldwide limited warranty and technical support is available at:
<http://h18004.www1.hp.com/products/servers/platforms/warranty/index.html>

HP Services and Support

The HP Care Pack service for ProLiant BL c-Class and p-Class server blades cover the server blade and all HP branded hardware options qualified for the server, purchased at the same time or afterwards, internal to the server. HP Care Pack Services provide total care and support expertise with committed response designed to meet your IT and business needs.

To fully capitalize on the capabilities of your HP BladeSystem servers, a service partner are required who thoroughly understands your server technology and systems environment. HP Services, an industry leader in the provisioning of multi vender support solutions, provides a range of support services designed to meet the varying needs of today's businesses. Whether an SMB or large global corporation, HP has a BladeSystem server support offer to help you rapidly deploy and maximize system uptime.

Recommended Service - Simplify BladeSystem solution implementation, maintenance, and management.

- Support Plus24 - 3 year 4 hour response 24x7 same business day coverage.
- Deployment Service - Installation and start up for HP BladeSystem Infrastructure.

Enhanced Service - Optimum service level to increase IT performance and availability.

- Support - 1 year HP Proactive BladeSystem Service.
- Deployment Service - Enhanced Network Installation and start up for HP BladeSystem Switches.

Installation & Start-Up service for HP BladeSystem Infrastructure plus HP BladeSystem Enhanced Network Installation and Start-UP as per the Customer Description and/or Data Sheet. To be delivered on a scheduled basis 8am-5pm, M-F, excluding HP holidays.

For a complete listing of service offerings and information visit:

<http://www.hp.com/services/bladessystemservices>
<http://www.hp.com/go/proliant/carepack>

Technical Specifications

Compliance	Mezzanine:	Dual-port QDR HCA: PCI Express 2.0 (1.1 compliant) Dual-port DDR HCA: PCI Express 2.0 (1.1 compliant) ROHS-R5 (QLogic HCA is ROHS-R6 compliant)
	Switches:	IBTA version 1.2 compatible ROHS-R5 (QLogic switch blade is ROHS-R6 compliant)
General Specifications	Communications Processor	Mezzanine: Dual-port QDR HCA: Mellanox ConnectX-2 Dual-port DDR HCA: Mellanox ConnectX-2 QLogic QDR HCA: QLogic TrueScale QDR switch: MT48436A1-FCC-Q, InfiniScale IV DDR Gen2 switch: MT48436A1-FCC-D, InfiniScale IV QLogic QDR switch: QLogic TrueScale
	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 11.4W Dual-port DDR HCA: max 12.9W QLogic QDR HCA: max 8.4W Switch blades: QDR switch: 10.8A at 12V Max (130W) DDR Gen2 switch: 10 A at 12V max (120 W) QLogic QDR switch: 150W @ 12V QLogic Management Module: 1.5W @ 12V
	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