

Mellanox ConnectX-2 VPI Single-port and Dual-port QDR InfiniBand Host Channel Adapters

Product Guide (withdrawn product)

High-performance computing (HPC) solutions require high bandwidth, low latency components with CPU offloads to get the highest server efficiency and application productivity. The Mellanox ConnectX-2 VPI Single-port and Dual-port Quad Data Rate (QDR) InfiniBand host channel adapters (HCAs) deliver the I/O performance that meets these requirements. Data centers and cloud computing also require I/O services such as bandwidth, consolidation and unification, and flexibility, and the Mellanox HCAs support the necessary LAN and SAN traffic consolidation.

Figure 1 shows the Mellanox ConnectX-2 VPI Dual-port QDR InfiniBand host channel adapter.




Figure 1. Mellanox ConnectX-2 VPI Dual-port QDR InfiniBand host channel adapter

Did you know?

Mellanox ConnectX-2 VPI Single-port and Dual-port QDR InfiniBand host channel adapters make it possible for any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network by using a consolidated software stack. With auto-sense capability, each ConnectX-2 port can identify and operate on InfiniBand, Ethernet, or Data Center Bridging (DCB) fabrics. ConnectX-2 with Virtual Protocol Interconnect (VPI) simplifies I/O system design and makes it easier for IT managers to deploy an infrastructure that meets the challenges of a dynamic data center.

Part number information

Table 1 shows the part numbers and feature codes for the Mellanox ConnectX-2 VPI QDR InfiniBand HCAs.

Table 1. Ordering part numbers and feature codes

Part number	Feature code	Description
81Y1531*	5446	Mellanox ConnectX-2 VPI Single-port QSFP QDR IB/10GbE PCI-E 2.0 HCA
81Y1535*	5447	Mellanox ConnectX-2 VPI Dual-port QSFP QDR IB/10GbE PCI-E 2.0 HCA

* Withdrawn from marketing

The adapters support the transceivers and direct-attach copper (DAC) twin-ax cables listed in Table 2.

Table 2. Supported transceivers and DAC cables

Part number	Feature code	Description
59Y1920	3731	3m QLogic Optical QDR InfiniBand QSFP Cable
59Y1924	3732	10m QLogic Optical QDR InfiniBand QSFP Cable
59Y1928	3733	30m QLogic Optical QDR InfiniBand QSFP Cable
59Y1892	3725	0.5m QLogic Copper QDR InfiniBand QSFP 30AWG Cable
59Y1896	3726	1m QLogic Copper QDR InfiniBand QSFP 30AWG Cable
59Y1900	3727	3m QLogic Copper QDR InfiniBand QSFP 28AWG Cable
49Y0488	5989	3m Optical QDR InfiniBand QSFP Cable
49Y0491	5990	10m Optical QDR InfiniBand QSFP Cable
49Y0494	5991	30m Optical QDR InfiniBand QSFP Cable

Figure 2 shows the Mellanox ConnectX-2 VPI Single-port QDR InfiniBand host channel adapter.




Figure 2. Mellanox ConnectX-2 VPI Single-port QDR InfiniBand host channel adapter

Features and benefits

The Mellanox ConnectX-2 VPI Single-port and Dual-port QDR InfiniBand host channel adapters have the following features:

InfiniBand

ConnectX-2 delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. Efficient computing is achieved by offloading from the CPU routine activities, which makes more processor power available for the application. Network protocol processing and data movement impacts, such as InfiniBand RDMA and Send/Receive semantics, are completed in the adapter without CPU intervention. Graphical processing unit (GPU) communication acceleration provides additional efficiencies by eliminating unnecessary internal data copies, which significantly reduces application runtime. The ConnectX-2 advanced acceleration technology enables higher cluster efficiency and scalability of up to tens of thousands of nodes.

RDMA over converged Ethernet

ConnectX-2 utilizes the InfiniBand Trade Association's RDMA over Converged Ethernet (RoCE) technology to deliver similar low-latency and high-performance over Ethernet networks. Leveraging Data Center Bridging capabilities, RoCE provides efficient low latency RDMA services over Layer 2 Ethernet. The RoCE software stack maintains existing and future compatibility with bandwidth and latency sensitive applications. With link-level interoperability in the existing Ethernet infrastructure, network administrators can use existing data center fabric management solutions.

TCP/UDP/IP acceleration

Applications utilizing TCP/UDP/IP transport can achieve industry leading throughput over InfiniBand or 10 GbE adapters. The hardware-based stateless offload engines in ConnectX-2 reduce the CPU impact of IP packet transport, allowing more processor cycles to work on the application.

I/O virtualization

ConnectX-2 with Virtual Intelligent Queuing (Virtual-IQ) technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server. I/O virtualization with ConnectX-2 gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

Storage accelerated

A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can use InfiniBand RDMA for high-performance storage access. T11 compliant encapsulation (FCoIB or FCoE) with full hardware offload simplifies the storage network while keeping existing Fibre Channel targets.

Software support

All Mellanox adapter cards are supported by a full suite of drivers for Microsoft Windows, Linux distributions, VMware, and Citrix XENServer. ConnectX-2 VPI adapters support OpenFabrics-based RDMA protocols and software. Stateless offload is fully interoperable with standard TCP/UDP/IP stacks. ConnectX-2 VPI adapters are compatible with configuration and management tools from OEMs and operating system vendors.

Specifications

The adapters have the following specifications:

- Low-profile adapter form factor
- Ports: One or two 40 Gbps InfiniBand interfaces (40/20/10 Gbps auto-negotiation) with QSFP connectors
- ASIC: Mellanox ConnectX-2
- Host interface: PCI Express 2.0 x8 (5.0 GTps)
- Interoperable with InfiniBand or 10G Ethernet switches

InfiniBand specifications:

- IBTA Specification 1.2.1 compliant
- RDMA, Send/Receive semantics
- Hardware-based congestion control
- 16 million I/O channels
- 256 to 4 KB MTU, 1 GB messages
- Nine virtual lanes: Eight data and one management

Enhanced InfiniBand specifications:

- Hardware-based reliable transport
- Hardware-based reliable multicast
- Extended Reliable Connected transport
- Enhanced Atomic operations
- Fine grained end-to-end quality of service (QoS)

Ethernet specifications:

- IEEE 802.3ae 10Gb Ethernet
- IEEE 802.3ad Link Aggregation and Failover
- IEEE 802.1Q, 1p VLAN tags and priority
- IEEE P802.1au D2.0 Congestion Notification
- IEEE P802.1az D0.2 ETS
- IEEE P802.1bb D1.0 Priority-based Flow Control
- Multicast
- Jumbo frame support (10 KB)
- 128 MAC/VLAN addresses per port

Hardware-based I/O virtualization:

- Address translation and protection
- Multiple queues per virtual machine
- VMware NetQueue support

Additional CPU offloads:

- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence
- Compliant with Microsoft RSS and NetDMA

Storage support:

- Fibre Channel over InfiniBand ready
- Fibre Channel over Ethernet ready

Management and tools:

InfiniBand:

- OpenSM
- Interoperable with third-party subnet managers
- Firmware and debug tools (MFT and IBDIAG)

Ethernet:

- MIB, MIB-II, MIB-II Extensions, RMON, and RMON 2
- Configuration and diagnostic tools

Protocol support:

- Open MPI, OSU MVAPICH, Intel MPI, MS MPI, and Platform MPI
- TCP/UDP, EoIB, IPoIB, SDP, and RDS
- SRP, iSER, NFS RDMA, FCoIB, and FCoE
- uDAPL

Physical specifications

The adapters have the following physical specifications (without the bracket):

- Single port: 2.1 in x 5.6 in (54 mm x 142 mm)
- Dual port: 2.7 in. x 6.6 in. (69 mm x 168 mm)

Operating environment

The adapters are supported in the following environment:

Operating temperature: 0 to 55° C

Air flow: 200 LFM at 55° C

Power consumption (typical):

- Single-port adapter: 7.0 W typical
- Dual-port adapter: 8.8 W typical (both ports active)

Power consumption (maximum):

- Single-port adapter: 7.7W maximum for passive cables only; 9.7W maximum for active optic modules
- Dual-port adapter: 9.4W maximum for passive cables only, 13.4W maximum for active optic modules

Warranty

One year limited warranty. When installed in an System x server, these cards assume your system's base warranty and any warranty upgrades.

Supported servers

The adapters are supported in the System x servers listed in Table 3.

Table 3. Server compatibility, part 1 (M5 systems and M4 systems with v2 processors)

Part number	Description	M5 systems (v3 processors)					M4 and X6 systems (v2 processors)											
		x3100 M5 (5457)	x3250 M5 (5458)	x3550 M5 (5463)	x3650 M5 (5462)	nx360 M5 (5465)	x3500 M4 (7383, E5-2600 v2)	x3530 M4 (7160, E5-2400 v2)	x3550 M4 (7914, E5-2600 v2)	x3630 M4 (7158, E5-2400 v2)	x3650 M4 (7915, E5-2600 v2)	x3650 M4 BD (5466)	x3650 M4 HD (5460)	x3750 M4 (8752, E5-4600 v2)	x3850 X6/x3950 X6 (3837)	dx360 M4 (7912, E5-2600 v2)	nx360 M4 (5455)	
81Y1531	Mellanox ConnectX-2 VPI Single-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
81Y1535	Mellanox ConnectX-2 VPI Dual-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 3. Server compatibility, part 2 (M4 systems with v1 processors and M3 systems)

Part number	Description	M4 and X5 systems (v1 processors)											M3 systems										
		x3100 M4 (2582)	x3250 M4 (2583)	x3300 M4 (7382)	x3500 M4 (7383, E5-2600)	x3530 M4 (7160)	x3550 M4 (7914, E5-2600)	x3630 M4 (7158)	x3650 M4 (7915, E5-2600)	x3690 X5 (7147)	x3750 M4 (8722)	x3850 X5 (7143)	dx360 M4 (7912, E5-2600)	x3200 M3 (7327, 7328)	x3250 M3 (4251, 4252)	x3400 M3 (7378, 7379)	x3500 M3 (7380)	x3550 M3 (7944)	x3620 M3 (7376)	x3630 M3 (7377)	x3650 M3 (7945)	x3755 M3 (7164)	dx360 M3 (6391)
81Y1531	Mellanox ConnectX-2 VPI Single-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	N	N	N	N	N	N	N	Y	N	Y	N	N	N	N	N	Y	N	N	N	N	Y	N
81Y1535	Mellanox ConnectX-2 VPI Dual-port QSFP QDR IB/10GbE PCI-E 2.0 HCA	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Supported operating systems

The adapters support the following operating systems:

- SUSE Linux Enterprise Server (SLES) 10 and 11
- Red Hat Enterprise Linux (RHEL) 4, 5.3, 5.4
- Microsoft Windows Server 2003
- Microsoft Compute Cluster Server 2003
- Microsoft Windows Server 2008
- Microsoft Windows HPC Server 2008
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)
- VMware ESX Server 3.5/vSphere 4.0

Related publications

For more information, refer to these documents:

- Mellanox ConnectX-2 VPI Single-port and Dual-port QDR InfiniBand HCA product page:
<http://www.mellanox.com>
(Select **Products** > **InfiniBand cards** > **ConnectX-2** .)

Related product families

Product families related to this document are the following:

- [InfiniBand & Omni-Path Adapters](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo 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:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO 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 jurisdictions 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo 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 Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result 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 measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, TIPS0778, was created or updated on November 5, 2014.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/TIPS0778>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS0778>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
System x®
X5

The following terms are trademarks of other companies:

Intel® is a trademark of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

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

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