

ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter

Product Guide

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter is a high-performance 25Gb Ethernet adapter suitable for Lenovo Flex System servers. The adapter is for customers who looking for end-to-end 25 Gb Ethernet speeds in their Flex System environment as well as those who want to maintain their existing 10Gb networking infrastructure while preparing for future upgrades to 25Gb network speeds.

The adapter has the standard form factor shape for Flex System mezzanine adapters.



Figure 1. ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter

Did you know?

The ConnectX-4 Lx adapter continues a long tradition of technology available for the Flex System ecosystem. The adapter enables customers to extend their investment in Flex System by offering high-speed connectivity into their networking infrastructure.

Part number information

The following table lists the ordering information for the ConnectX-4 Lx adapter.

Table 1. Part number information

Part number	Feature code	Description
4XC7A08315	BHBX	ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter

Features

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter has the following key features:

- **I/O Virtualization**
ConnectX-4 Lx EN SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VMs) within the server. I/O virtualization with ConnectX-4 Lx EN gives data center administrators better server utilization while reducing cost, power, and cable complexity, allowing more virtual machines and more tenants on the same hardware.
- **Overlay Networks**
In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX-4 Lx EN effectively addresses this by providing advanced NVGRE, VXLAN and GENEVE hardware offloading engines that encapsulate and de-capsulate the overlay protocol headers, enabling the traditional offloads to be performed on the encapsulated traffic for these and other tunneling protocols (GENEVE, MPLS, QinQ, and so on). With ConnectX-4 Lx EN, data center operators can achieve native performance in the new network architecture.
- **RDMA over Converged Ethernet (RoCE)**
ConnectX-4 Lx EN supports RoCE specifications delivering low-latency and high-performance over Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as ConnectX-4 Lx EN advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.
- **Mellanox PeerDirect**
Mellanox PeerDirect communication provides high efficiency RDMA access by eliminating unnecessary internal data copies between components on the PCIe bus (for example, from GPU to CPU), and therefore significantly reduces application run time. ConnectX-4 Lx EN advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.
- **Storage Acceleration**
Storage applications will see improved performance with the higher bandwidth ConnectX-4 Lx EN delivers. Moreover, standard block and file access protocols can leverage RoCE for high-performance storage access. A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks.

Specifications

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter has the following technical specifications:

- Mellanox ConnectX-4-Lx ASIC
- PCIe 3.0 x16 host interface (standard Flex System compute node I/O adapter connection)
- Two copper 25GBase-KR SERDES interfaces (2 lanes of 25G each)
- Connections to the I/O module bays in the Flex System Enterprise Chassis
 - Mezz slot 1 to I/O module bays 1 and 2
 - Mezz slot 2 to I/O module bays 3 and 4
 - All connections are internal to the Flex chassis; no transceivers or cables are required
- Supports Message Signal Interrupt (MSI-X)
- Support for PXE boot, iSCSI boot and Wake-on-LAN (WOL)
- Networking Features
 - Jumbo frames (up to 9600-Byte)
 - 802.3x flow control
 - Link Aggregation (IEEE 802.1AX-2008)
 - Virtual LANs-802.1q VLAN tagging
 - Configurable Flow Acceleration
 - Congestion Avoidance
- CPU Offloads
 - RDMA over Converged Ethernet (RoCE)
 - TCP/UDP/IP stateless offload
 - LSO, LRO, checksum offload
 - RSS (can be done on encapsulated packet), TSS, VLAN insertion/ stripping, Receive flow steering
 - Intelligent interrupt coalescence
- Overlay Networks
 - Stateless offloads for overlay
 - networks and tunneling protocols
 - Hardware offload of encapsulation and decapsulation of NVGRE and VXLAN overlay networks
- Hardware-Based I/O Virtualization
 - Single Root IOV
 - Multi-function per port
 - Address translation and protection
 - Multiple queues per virtual machine
 - Enhanced QoS for vNICs
 - VMware NetQueue support
- Virtualization
 - SR-IOV: Up to 256 Virtual Functions
 - SR-IOV: Up to 8 Physical Functions per port
 - Virtualization hierarchies (e.g. NPAR)
 - Virtualizing Physical Functions on a physical port
 - SR-IOV on every Physical Function
 - Ingress and egress QoS levels
 - Guaranteed QoS for VMs
- Protocol Support
 - OpenMPI, IBM PE, OSU MPI (MVAPICH/2), Intel MPI
 - Platform MPI, UPC, Open SHMEM
 - TCP/UDP, MPLS, VxLAN, NVGRE,

- o GENEVE
- o iSER, NFS RDMA, SMB Direct
- o uDAPL

IEEE Standards

The adapter supports these IEEE specifications:

- IEEE 802.3by, 25 Gb/s supporting all FEC modes
- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE 802.3az Energy Efficient Ethernet (supports only “Fast-Wake” mode)
- IEEE 802.3ap based auto-negotiation and KR startup
- IEEE 802.3ad, 802.1AX Link Aggregation
- IEEE 802.1Q, 802.1P VLAN tags and priority
- IEEE 802.1Qau (QCN) Congestion Notification
- IEEE 802.1Qaz (ETS)
- IEEE 802.1Qbb (PFC)
- IEEE 802.1Qbg
- IEEE 1588v2
- PCI Express Gen 3.0

The adapter supports these additional specifications:

- IPv4 (RFQ 791)
- IPv6 (RFC 2460)

Server support

The following table lists the compute nodes that support the adapters.

Table 2. Support for Flex System compute nodes

Part number	Description	x240 (8737, E5-2600 v2)	x240 (7162)	x240 M5 (9532, E5-2600 v3)	x240 M5 (9532, E5-2600 v4)	x440 (7167)	x880/x480/x280 X6 (7903)	x280/x480/x880 X6 (7196)	SN550 (7X16)	SN850 (7X15)	SN550 V2 (7Z69)
4XC7A08315	ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter	N	N	N	N	N	N	N	Y	Y	Y

I/O adapter cards are installed in the slot in supported servers, such as the SN550 V2, as highlighted in the following figure.




Figure 2. Location of the I/O adapter slots in the SN550 V2

Internal connectivity

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter does not use transceivers or cables for connectivity to switches. Instead, the adapter connects to the Flex System switches and I/O modules installed in the chassis via internal connections.

The following table shows the connections between adapters installed in the compute nodes and the switch bays in the chassis.

Table 3. Adapter to I/O bay correspondence

I/O adapter slot in the compute node	Port on the adapter	Corresponding I/O module bay in the chassis			
		Bay 1	Bay 2	Bay 3	Bay 4
Slot 1	Port 1	Yes			
	Port 2		Yes		
Slot 2	Port 1			Yes	
	Port 2				Yes
Slot 3 (SN850 only)	Port 1	Yes			
	Port 2		Yes		
Slot 4 (SN850 only)	Port 1			Yes	
	Port 2				Yes

The connections between the adapters installed in the compute nodes to the switch bays in the chassis are shown diagrammatically in the following figure. The figure shows half-wide servers (such as the SN550 V2 with two adapters) and full-wide servers (such as the SN850 with four adapters).

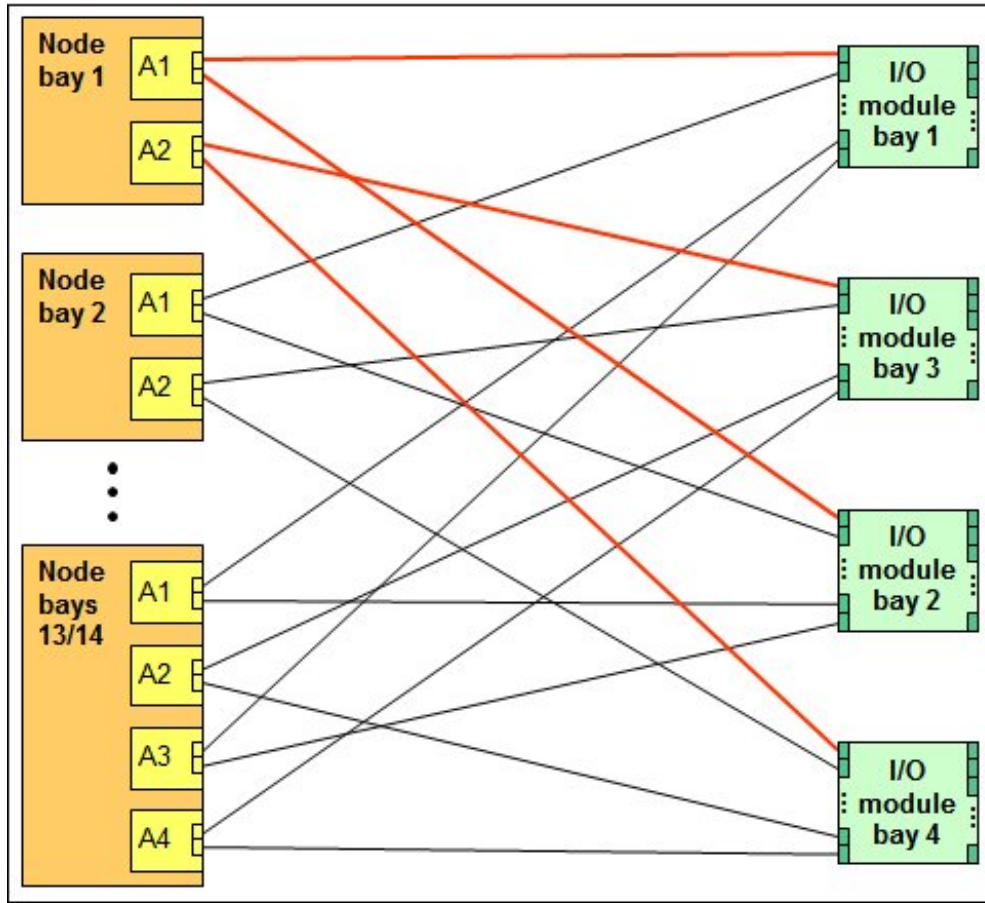


Figure 3. Logical layout of the interconnects between adapters and I/O modules

I/O module support

These adapters can be installed in any I/O adapter slot of a supported Flex System compute node. One or two compatible modules must be installed in the corresponding I/O bays in the chassis.

The following table lists the supported I/O modules.

Table 4. I/O modules supported

Part number	Description
4SG7A08868	Lenovo ThinkSystem NE2552E Flex Switch

Operating system support

The following table lists the supported operating systems.

Tip: This table is automatically generated based on data from [Lenovo ServerProven](#).

Table 5. Operating system support for ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter, 4XC7A08315

Operating systems	SN550 V2	SN550 (Xeon Gen 2)	SN850 (Xeon Gen 2)	SN550 (Xeon Gen 1)	SN850 (Xeon Gen 1)
Microsoft Windows Server 2016	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.6	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.7	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.8	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.9	Y	Y	Y	N	N
Red Hat Enterprise Linux 9.0	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 9.1	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 9.2	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 9.3	Y	Y	Y	N	N
Red Hat Enterprise Linux 9.4	Y	N	N	N	N
SUSE Linux Enterprise Server 12 SP4	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y

	SN550 V2	SN550 (Xeon Gen 2)	SN850 (Xeon Gen 2)	SN550 (Xeon Gen 1)	SN850 (Xeon Gen 1)
Operating systems					
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP4	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP5	Y	Y	Y	Y	Y
Ubuntu 22.04 LTS	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 8.0	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 8.0 U1	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 8.0 U2	Y	Y	Y	Y	Y

Warranty

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter has a 1-year, customer-replaceable unit (CRU) limited warranty. When installed in a server, the adapter assumes the server's base warranty and any warranty upgrade purchased for the server.

Physical specifications

The adapters have the following dimensions and weight:

- Width: 100 mm (3.9 in.)
- Depth: 80 mm (3.1 in.)
- Weight: 13 g (0.3 lb)

The adapters have the following shipping dimensions and weight (approximate):

- Height: 58 mm (2.3 in.)
- Width: 229 mm (9.0 in.)
- Depth: 208 mm (8.2 in.)
- Weight: 0.4 kg (0.89 lb)

Regulatory compliance

The ThinkSystem Mellanox ConnectX-4 Lx 25Gb 2-port Mezz Adapter conforms to the following regulatory standards:

- Safety: CB, cTUVus, CE
- EMC: CE, FCC, VCCI, ICES, RCM
- RoHS: RoHS-R6

Related publications and links

For more information, see the following resources:

- ThinkSystem NE2552E Flex Switch Product Guide
<http://lenovopress.com/LP0854>
- ServerProven
<http://www.lenovo.com/us/en/serverproven>

Related product families

Product families related to this document are the following:

- [25 Gb Embedded Connectivity](#)
- [Blade Network 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, LP1559, was created or updated on May 11, 2022.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1559>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1559>.

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®

Flex System

ServerProven®

ThinkSystem®

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks 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.