



LinkX User Guide for 400G and 200G using 50G-PAM4 and 100G using 25G-NRZ Modulation Cables and Transceivers

Table of Contents

- NVIDIA LinkX QSFP-based Cables and Transceiver Portfolio 3**
 - LinkX 400GbE Cables and Transceivers Portfolio 3
 - LinkX 200GbE and InfiniBand HDR Cables and Transceiver Portfolio 5
 - NVIDIA Spectrum 400G, 200G, 100GbE Ethernet Switch Portfolio 7
 - NVIDIA Quantum HDR 200G InfiniBand Switch Portfolio 8
 - NVIDIA ConnectX Network Adapter Portfolio 8
 - NVIDIA BlueField-3 DPU Portfolio 10
- Where to Find More LinkX Documentation 11**
- Document Revision History 12**

NVIDIA LinkX QSFP-based Cables and Transceiver Portfolio

LinkX 400GbE Cables and Transceivers Portfolio

This document profiles 400Gb/s QSFP-DD-based cables and transceivers for Ethernet only from 0.5-meters to 10km

The LinkX® 400GbE QSFP-DD portfolio consists of straight and splitter DACs, an AOC cable and five transceivers with reaches from 1-to 10k-meters. QSFP-DD cables and transceivers support only Ethernet.

NVIDIA 400GbE QSFP-DD product line:

- DAC cables have 1:2 and 1:4 split configuration with maximum reach of 2.5-meters
- AOC have only QSFP-DD-to-QSFP-DD AOC and no splitters.
- 400GbE SR8 multimode transceiver can be split to 2x 200G QSFP56 and to 2x 100G
 - Use 3rd party MPO-16/APC-to-2x MPO-12/UPC fiber cables - not offered by NVIDIA.
 - Note, the 400GbE SR8 uses a 16-fiber MPO-16/APC optical connector
- 400GbE DR4 can be split to 4x 100G DR1 using 3rd party supplied splitter fiber cables.
 - MPO-12/APC-to-4x LC Duplex
 - DR1 is 1x100G-PAM4 optical to 4x25G-NRZ
- 400G FR4 and LR4 are optically multiplexed and cannot be split and used for long, leaf-spin and spine-super spine links.

LinkX 400GbE Ethernet QSFP-DD Cables and Transceiver Portfolio

400GBE ETHERNET QSFP-DD (50G-PAM4)

SN5400
400GbE Spine/Super Spine Ethernet Switch
64 400GbE QSFP-DD cages
Spectrum-4





SN4700
400GbE Spine/ToR Ethernet Switch
32 QSFP-DD cages
Spectrum-3



Direct Attach Copper (DAC)



400GbE
MCP1660-W0xxExx



400GbE-to-2x200GbE
MCP7H60-W0xxRxx



400GbE-to-4x100GbE
MCP7F60-W0xxRxx

Active Optical Cables (AOCs)

400GbE
C-DQ8FNMxxx-H0-M



Switch-to-Switch or
ConnectX-6 and/or
BlueField-2 DPUs

400GbE SR8 QSFP-DD 100m

T-DQ8FNS-N00-M



MPO-16/APC



400GbE DR4 QSFP-DD 500m

MMS1V00-WM



MPO-12/APC



400GbE FR4 QSFP-DD 2km

MMS1V50-WM




400GbE LR4 QSFP-DD 10km

MMS1V90-WR

Duplex LC/UPC
2-fiber
Optical Connecto




SN5400
400GbE Spine/Super Spine Ethernet Switch





SN4700
400GbE Spine/ToR Ethernet Switch



LinkX 200GbE and InfiniBand HDR Cables and Transceiver Portfolio

The LinkX 200GbE and InfiniBand HDR QSFP56 portfolio consists of straight and 1:2 splitter DACs and AOC cables, two transceivers with reaches from 1-to 2k-meters. Some of the parts are either InfiniBand or Ethernet specific with different part numbers, or sometimes a single part number supports both InfiniBand and Ethernet.

- 200GbE and HDR portfolios are very similar, except the InfiniBand DACs are maximum reach of 2-meters.
- 200G straight DAC and 200G SR4 multimode transceivers are InfiniBand part number specific and require more stringent BER testing.
- InfiniBand cables and transceivers can be used in NVIDIA Spectrum™ Ethernet switches in NVIDIA end-to-end configurations. InfiniBand testing is more stringent than Ethernet. However, Ethernet cables and transceivers cannot be used in InfiniBand systems.
- 200GbE and HDR SR4 transceivers can be split into 2x 100G creating 2x50G-PAM4 QSFP56 100GbE or HDR100 ends.
 - Use 3rd party MPO-12/UPC-to-MPO-12/UPC fiber cables -- not offered by NVIDIA.
- 200G FR4 is optically multiplexed and cannot be split and used for long, leaf-spin and spine-super spine links.
- 200G FR4 can be used in ConnectX®-6 adapters under specific air flow conditions in servers; check the *NVIDIA ConnectX-6 User Manual*.
- 200GbE and HDR use 4x50G PAM4 = 200G.
- HDR100 and 100GbE based on 2x50G PAM4 is 100G and not compatible with 100GbE based on 4x25G-NRZ.

LinkX 200GbE Ethernet & HDR InfiniBand QSFP56 Cables and Transceiver Portfolio

200GBE & HDR INFINIBAND QSFP56 (50G-PAM4)



200G QSFP56-to-200G QSFP56

MCP1650-H0xxExx IB only
MCP1650-H0xxExx EN only



MFS1500-H0xxV



IB+ EN

200G QSFP56-to-2x 100G QSFP56

MCP7H50-H0xxRxx IB+ EN



200G QSFP56-to-4x 50G SFP56

MCP7H70-V0xxRxx EN only



MFS1550-H0xxV IB+ EN



200Gb SR4 HDR 850nm Multi-mode 100-meters QSFP56

MMA1T00-HS IB-only
MMA1T00-VS EN only



200Gb FR4 HDR 1310nm Single-mode 2km QSFP56

MMS1W50-HM IB+ EN



Fibers not supplied by NVIDIA



NVIDIA Spectrum 400G, 200G, 100GbE Ethernet Switch Portfolio

There are many different switch port configurations in the Spectrum Ethernet switch product line using various combination of QSFP-DD, QSFP56, QSFP28 and SFP28. The DACs, AOCs and multimode transceiver are generally used in switch-to server links and multimode for longer links beyond 2.5-meter DAC maximum reach.

The NVIDIA Spectrum, Spectrum-2, Spectrum-3, and Spectrum-4 are successive switch IC generations dating back to 2015 with the Spectrum-4 being the latest offering. The Spectrum-4 is also offered as the SN5600 in twin-port **OSFP** 800G 64-port configuration based on 100G-PAM4 modulation. The Spectrum Ethernet switches are protocol specific and do not support InfiniBand.

NVIDIA Spectrum Ethernet Switch Portfolio

SN2000 Series Spectrum QSFP28 & SFP28	SN3000 Series Spectrum-2 QSFP56, QSFP28 & SFP28	SN4000 Series Spectrum-3 QSFP-DD, QSFP56 & SFP28	SN5000 Series Spectrum-4 QSFP-DD, QSFP56
			

The 400GbE QSFP-DD cables and transceivers line is for the NVIDIA SN5400, SN4700, SN4410 Ethernet switches only and not used for InfiniBand. The QSFP-DD cages are backwards compatible and accept all QSFP-based cables and transceivers, both 50G-PAM4 and 25G-NRZ line rates and 25, 40, 100, 200, and 400G aggregate data rates. Similarly, the 200GbE QSFP56 cages are backwards compatible to QSFP28.

<ul style="list-style-type: none"> • 400GbE • 200GbE • 100GbE 	QSFP-DD QSFP56 QSFP28	8x50G-PAM4 4x50G-PAM4 4x25G-NRZ
--	-----------------------------	---------------------------------------

NVIDIA Quantum HDR 200G InfiniBand Switch Portfolio

The 200G HDR is offered in 32-ports of 200G HDR in QSFP56 cages.

- QSFP56 cages accept 100G EDR cables and transceivers
- HDR was the first InfiniBand products to offer 1:2 splitter DACs and AOCs

The HDR cables and transceivers product line overlaps the 200GbE Ethernet portfolio with a few InfiniBand-only exceptions for a DAC and SR4 transceiver. InfiniBand has a much lower bit error ratio (BER) requirement to minimize the use of forward error correction which induces significant latency.

NVIDIA Quantum HDR InfiniBand Switch Portfolio
 NVIDIA Quantum QM8700 Switch
 32 cages QSFP56 200Gb/s



NVIDIA ConnectX Network Adapter Portfolio

The ConnectX line converts PCIe bus signals to Ethernet or InfiniBand protocols. The ConnectX portfolio consists of three generation of adapter ICs.

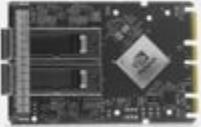
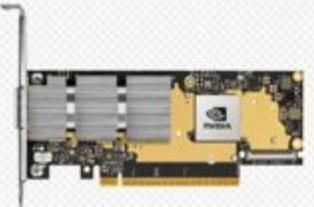
<ul style="list-style-type: none"> • 400GbE/NDR • 200GbE/HDR • 100GbE/EDR 	ConnectX-7 ConnectX-6 ConnectX-5	QSFP112 QSFP56 QSFP28	OSFP SFP56 SFP28
--	--	-----------------------------	------------------------

- ConnectX-7 is offered in 4x100G-PAM4 OSFP and QSFP112.
- ConnectX-7 and most ConnectX-6 products support both Ethernet and InfiniBand.
- SFP56 and OCP cards are Ethernet specific.
- As with the switch port cages, the ConnectX-7 QSFP112 is backwards compatible to the slower QSFP56, QSFP28, etc. Similarly, for the QSFP56 based ConnectX-6 and QSFP28.
- ConnectX line is not offered with QSFP-DD cages.

NVIDIA ConnectX-7 Network Adapter Portfolio

CONNECTX-7 PORTFOLIO

Comprehensive portfolio to support entry level & premium boards

<p>4P X SFP56 - HHHL / FHHL 25G/50G - W/ & W/O SYNCE/GNSS</p>	<p>2P X QSFP112 - HHHL 200G - PAM4 56 & PAM4 112</p>	<p>2P X QSFP112 - OCP3.0 200G - PAM4 56 & PAM4 112</p>
		
<p>1P X OSFP - HHHL 400G - PAM4 112</p>	<p>1P X OSFP - OCP3.0 400G - PAM4 112</p>	
		

NVIDIA BlueField-3 DPU Portfolio

BlueField DPUs contain an ARM® CPU, memory, accelerators, PCIe switch and ConnectX-6 or ConnectX-7 adapters.

- BlueField-2 DPUs are offered in 200G QSFP56 for InfiniBand and Ethernet
- BlueField-3 DPUs are offered in 200G and 400G QSFP112 for InfiniBand and Ethernet
- Not offered with OSFP or QSFP-DDs
- QSFP112 cages accept 200G QSFP56 and 100G QSFP28

NVIDIA BlueField-3 DPU Portfolio 200G & 400G, InfiniBand and Ethernet QSFP112

Single QSFP112	Dual QSFP112
 A photograph of a single NVIDIA BlueField-3 DPU. The board is populated with various components, including a central chip with the NVIDIA logo, memory modules, and a single QSFP112 port on the left side.	 A photograph of a dual NVIDIA BlueField-3 DPU. The board is populated with various components, including a central chip with the NVIDIA logo, memory modules, and two QSFP112 ports on the left side.

Where to Find More LinkX Documentation

This user guide is to be used in conjunction with other documents located in folders in docs.nvidia.com/networking/ > Interconnect. This site is where the following LinkX cables and transceivers documents are provided.

LinkX Overview Documents:	Review of parts, important notes, and configuration details for linking to NVIDIA switches and adapters <ul style="list-style-type: none">• LinkX Cables and Transceivers Guide to Key Technologies• LinkX User Guide for 400Gb/s 100G-PAM4 OSFP & QSFP112-based Cables and Transceivers• LinkX User Guide for 400Gb/s and 200Gb/s using 50G-PAM4 and 100Gb/s using 25G-NRZ Modulation Cables and Transceivers (this document)
Configuration Maps:	Picture and part number-based PowerPoint® slides for every configuration with NVIDIA switches, network adapters, and DGX GPU systems for 100G-PAM4, 50G-PAM4, 25G-NRZ cables and transceivers <ul style="list-style-type: none">• Configuration Maps
Parts Lists:	Tables summarize by speed, form factor, connector, power, reach, etc. and hyperlinks to individual products specs <ul style="list-style-type: none">• 400Gb/s (100G-PAM4) Transceivers and Fiber Parts List• 400Gb/s and 200Gb/s (50G-PAM4) and 100Gb/s (25G-NRZ) Cables and Transceivers Parts List using QSFP-DD, QSFP56, QSFP28, SFP28
Product Specifications:	10-to-16-page detailed hardware datasheets with physical, thermal, electrical, and optical specifications for each product <ul style="list-style-type: none">• docs.nvidia.com/networking/ > Interconnect > <i>select speed and type</i>
Additional Docs:	<ul style="list-style-type: none">• NVIDIA Cable Management Guidelines and FAQ

Document Revision History

Version	Date	Changes
1.0	August 2023	Initial release

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete. NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT,



INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA Corporation & affiliates. All Rights Reserved.

