NVIDIA MMA1Z00-NS400 400Gb/s, Single-port, QSFP112 Multimode SR4 Transceivers Product Specifications

Table of Contents

1 Introduction 3
1.1 Key Features 3
1.2 Applications 4
2 Overview 5
2.1 Transceiver Connectivity Scenarios 5
2.2 Use cases 5
3 Pin Description 7
3.1 QSFP112 Pin Description 7
3.1.1 Pin definitions of the module high speed inputs/outputs: 8
3.2 Control Signals 8
3.3 Diagnostics and Other Features 9
4 Specifications 10
4.1 Absolute Maximum Specifications 10
4.2 Recommended Operating Conditions and Power Supply Requirements 10
4.3 Electrical Specifications 11
4.3.1 Electrical Specification for Low Speed Signal 11
4.4 Optical Specifications 11
4.4.1 New Form-factors 14
4.5 Labels 14
4.5.1 Back shell Label 14
4.5.1.1 Transceiver Back-Shell Label Serial Number Legend 15
4.5.2 Regulatory Compliance 15
4.6 Connector and Cabling Details 16
4.6.1 MPO-12/APC Optical Connector 16
4.6.1.1 NVIDIA Supplied Crossover Type-B Fiber Cables 17
4.6.2 Handling and Cleaning 20
4.6.3 Cable Management Guidelines 20
5 Ordering Information 21
5.1 Part Numbers and Description 21
6 Recommended NVIDIA Supplied Crossover Fiber Cables Part Numbers 22
7 Document Revision History 24

1 Introduction

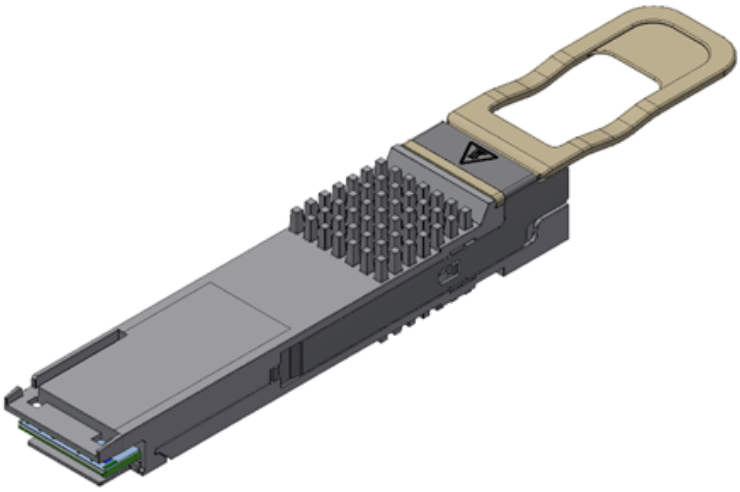
The NVIDIA MMA1Z00-NS400 is an InfiniBand and Ethernet 400Gb/s, Single-port, QSFP112, SR4 multimode parallel transceiver using a single, 4-channel MPO-12/APC optical connector. The Short Reach 4-channel (SR4) design uses 100G-PAM4 modulation and has a maximum fiber reach of 50-meters using OM4 multimode fiber and assumes two optical patch panels in the link. It has identical design and internals as the OSFP version, only with different connector shells.

The transceiver firmware supports both InfiniBand and Ethernet and is automatically enabled depending on the protocol of the switch attached to. The QSFP112 shell has a flat-top and utilizes the riding heat sink (cooling fins) on the ConnectX-7 or BlueField-3 connector cage. The small bumps near the pull tab provide additional cooling and remains outside the host connector cage.

When linked to 1:2 splitter fiber cable split end has only 2 channels and will activate only 2-channels in the 400G transceiver automatically creating a 200G speed and reducing power.

Multimode optics is denoted by a tan-colored pull tab and aqua-colored optical fiber. Green plastic shell on the MPO-12/APC optical connector denotes Angled Polish Connector (APC) and is not compatible with the aqua colored Ultra-flat Polished Connectors (UPC).

NVIDIA's Single-port and Twin-port transceiver combinations guarantee optimal operation in NVIDIA end-to-end InfiniBand systems and a rigorous production tested to ensure the best out-of-the-box installation experience, performance, and durability.



Images are for illustration purposes only. Product labels, colors, and lengths may vary.

1.1 Key Features

- IB and ETH support
- 400G SR4 multimode
- 4-channels of 100G-PAM4 modulation
- QSFP112 connector shell
- 850nm wavelength VCSEL laser
- Single MPO-12/APC optical connector

- Max reach:
    - 30m on OM3
    - 50m on OM4
- 8.5W max (4-channels)
- 5.5W max (2-channels)
- Single 3.3V power supply
- Class 1 laser safety
- Hot pluggable, RoHS based
- QSFP112 MSA
- CMIS 4.0 compliant
- Case temperature range of 0°C to +70°C

1.2 Applications

- Used in ConnectX-7/QSFP112 adapters or BlueField-3/QSFP112 DPUs linked to Twin-port transceivers in 2x400G IB/EN switches

2 Overview

2.1 Transceiver Connectivity Scenarios

The transceiver is used for connecting 400G and 200Gb/s BlueField-3/QSFP112 Data Processing Units (DPU) or ConnectX-7/QSFP112-based, PCIe-bus network cards. Typically, the transceiver is linked to a single 800Gb/s Twin-port 2x400G OSFP transceiver (MMA4Z00-NS) in a Quantum-2 InfiniBand or Spectrum-4 Ethernet switch. The 400Gb/s transceiver has two speeds depending on the number of fibers attached:

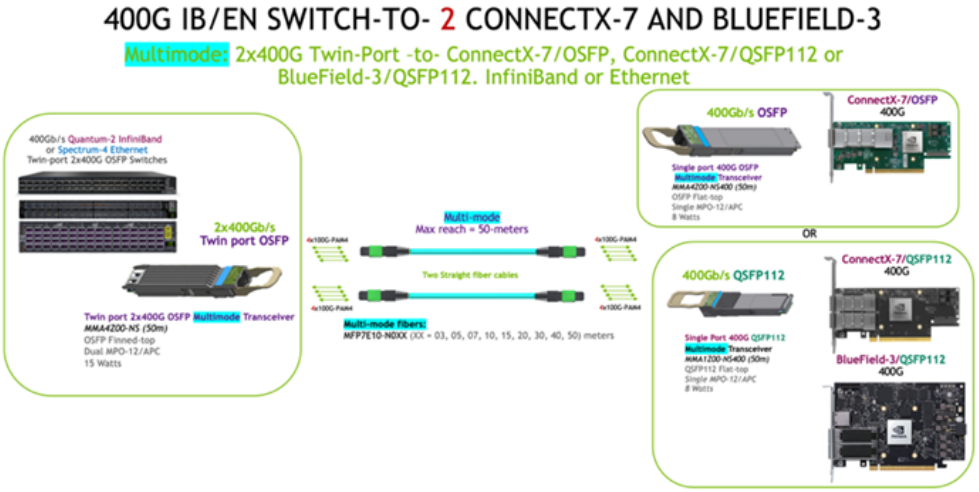
1. 400Gb/s mode: Using 4-channels and straight 50-meter crossover fiber cables (MFP7E10), the transceiver draws 9 Watts maximum or 8 Watts typical. In this case, the Twin-port 2x400G transceiver supports 400G transceivers in two ConnectX-7/QSFP112 and/or two BlueField-3/QSP112 DPU cards.
2. 200Gb/s mode: Using 2-channels and 1:2 splitter 50-meter crossover fiber cables (MFP7E20), the transceiver operates at 200Gb/s NDR200 rate and draws 5 Watts maximum. It automatically reduces power from 8 Watts as only 2 channels are activated. This case creates links to four 200Gb/s ConnectX-7/QSFP112 of BlueField-3/QSFP112 adapter cards.



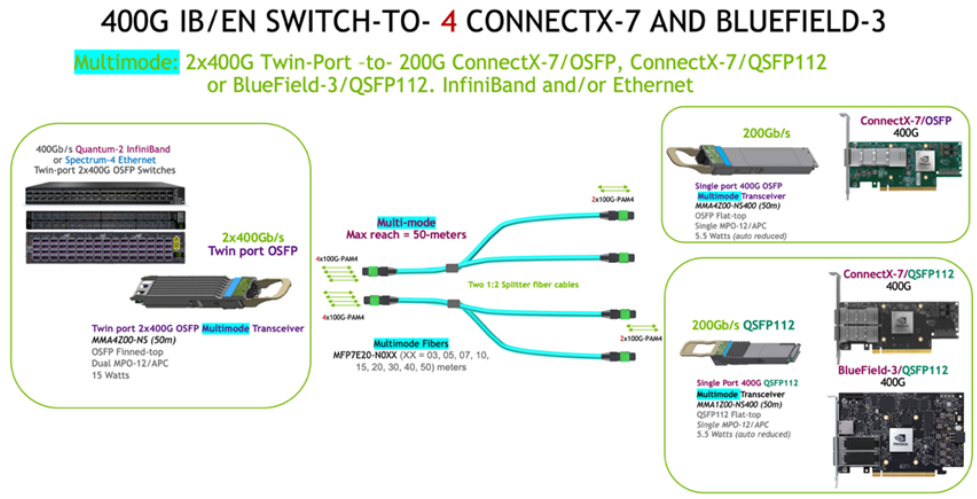
- QSFP112 are not for use in switches. BlueField-3 only accepts QSFP112s
- Both fibers in the Twin-port 2x400G transceiver linked to the QSFP112s must be the same type – straight or splitter and cannot be mixed

2.2 Use cases

1. Switch-to-two 400G ConnectX-7/QSFP112 or BlueField-3/QSFP112
   A Twin port OSFP transceiver using two, straight fiber cables can support up to two ConnectX-7/QSFP112 adapters and/or two BlueField-3 DPUs. Each of the two, 4-channel fiber cables (MFP7E10) can link to the 400G QSFP112 MMA1Z00-NS400 transceiver up to 50-meters.
   - ConnectX-7 adapters are offered on both OSFP and QSFP112
   - BlueField-3 adapters only accept QSFP112 devices



2. Switch-to-four 200G ConnectX-7/QSFP and/or BlueField-3/QSFP112
   A Twin port OSFP transceiver using two, 1:2 fiber splitter cables can support up any combination of four ConnectX-7 adapters and/or BlueField-3/QSFP1212 DPUs. Each of the two, 4-channel 1:2 fiber splitter cables (MFP7E20) can link to a 400G QSFP112 MMA1Z00-NS400 transceiver up to 50-meters.
   The two-fiber channel ends only activate two of the lanes in the 400G transceiver creating a 200G device and automatically reduces the power consumption of only the 400G transceivers from 8 Watts typical to 5.5 Watts typical. Twin port OSFP power consumption remains at 15 Watts.

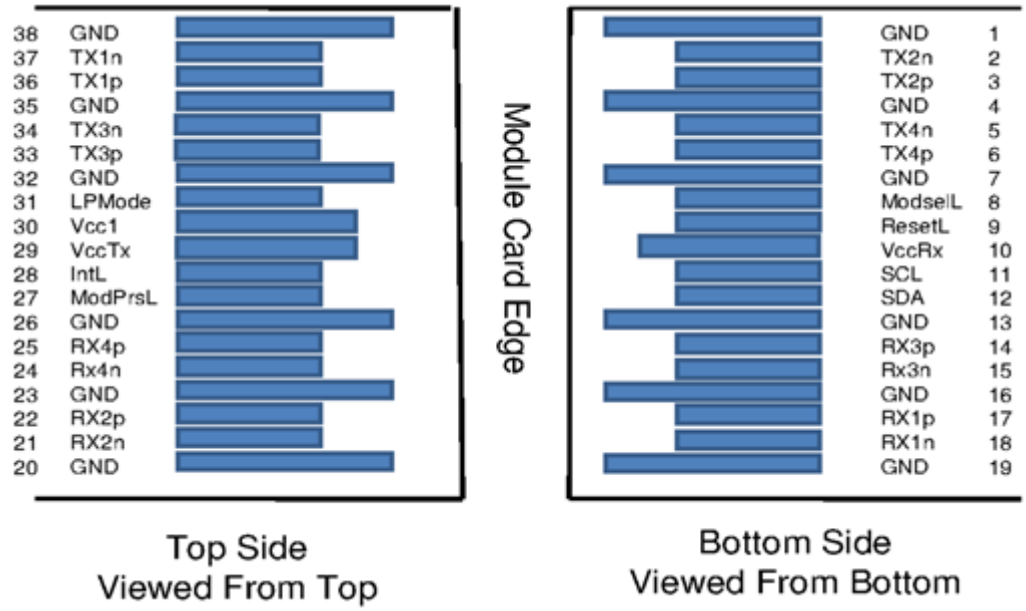


3 Pin Description

3.1 QSFP112 Pin Description

Pin	Symbol	Description	Pin	Symbol	Description
1	GND	Ground	20	GND	Ground
2	Tx2n	Transmitter Inverted Data Input	21	Rx2n	Receiver Inverted Data Output
3	Tx2p	Transmitter Non-Inverted Data Input	22	Rx2p	Receiver Non-Inverted Data Output
4	GND	Ground	23	GND	Ground
5	Tx4n	Transmitter Inverted Data Input	24	Rx4n	Receiver Inverted Data Output
6	Tx4p	Transmitter Non-Inverted Data Input	25	Rx4p	Receiver Non-Inverted Data Output
7	GND	Ground	26	GND	Ground
8	ModSelL	Module Select	27	ModPrsL	Module Present
9	ResetL	Module Reset	28	IntL	Interrupt
10	Vcc Rx	+3.3V Power Supply Receiver	29	Vcc Tx	+3.3V Power Supply Transmitter
11	SCL	2-wire Serial Interface Clock	30	Vcc1	+3.3V Power Supply
12	SDA	2-wire Serial Interface Data	31	LPMode	Low Power Mode
13	GND	GND	32	GND	Ground
14	Rx3p	Receiver Non-Inverted Data Output	33	Tx3p	Transmitter Non-Inverted Data Input
15	Rx3n	Receiver Inverted Data Output	34	Tx3n	Transmitter Inverted Data Input
16	GND	Ground	35	GND	Ground
17	Rx1p	Receiver Non-Inverted Data Output	36	Tx1p	Transmitter Non-Inverted Data Input
18	Rx1n	Receiver Inverted Data Output	37	Tx1n	Transmitter Inverted Data Input
19	GND	Ground	38	GND	Ground

3.1.1 Pin *definitions* of the module high speed inputs/ outputs:



3.2 Control Signals

This transceiver is CMIS 4.0 compliant. This means that the control signals shown in the pad layout support the following functions:

Name	Description
ModPrsL	Module Present pin, grounded inside the module. Terminated with pull-up in the host system. Asserted low when the transceiver is inserted, whereby the host detects the presence of the transceiver.
ModSelL	Module Select input pin, terminated high in the module. Only when held low by the host, the module responds to 2-wire serial communication commands. The ModSelL enables multiple modules to share a single 2-wire interface bus.
ResetL	Reset input pin, pulled high in the module. A low level on the ResetL pin for longer than the minimum pulse length (t_Reset_init) initiates a complete module reset, returning all user module settings to their default state. During reset the host shall disregard all status bits until the module indicates completion of the reset interrupt by asserting IntL signal low with the Data_Not_Ready bit negated. Note that on power up (including hot insertion) the module completes the reset interrupt without requiring a reset.
LPMode	Low Power Mode input, pulled up inside the module. The transceiver starts up in low-power mode, i.e. <1.5 W with the two-wire interface active. The host system can read the power class declaration from the transceiver and determine if it has enough power to enable the high-speed operation/high power mode of the transceiver. This can be done by asserting LPMode low or by use of the Power_over-ride and Power_set control bits (Address A0h, byte 93 bits 0,1).

IntL	Interrupt Low is an open-collector output, terminated high in the host system. A “Low” indicates a possible module operational fault or a status critical to the host system, e.g. temperature alarm. The host identifies the source of the interrupt using the 2-wire serial interface. The INTL pin is de-asserted “High” after completion of reset, when byte 2 bit 0 (Data Not Ready) is read with a value of ‘0’.

3.3 Diagnostics and Other Features

The transceiver supports the QSFP112 MSA specification and has the following key features:

Physical layer link optimization:

- Programmable Tx input equalization
- Programmable Rx output amplitude
- Programmable Rx output pre-emphasis

Digital Diagnostic Monitoring (DDM):

- Rx receive optical power monitor for each lane
- Tx transmit optical power monitor for each lane
- Tx bias current monitor for each lane
- Supply voltage monitor
- Transceiver case temperature monitor
- Warning and Alarm thresholds for each DDM function (not user changeable)

Other SFF-8636 functions and interrupt indications:

- Tx & Rx LOS indication
- Tx & Rx LOL indication
- Tx fault indication

LOS, LOL, and Tx Fault status flags can be read via the two-wire management interface and are jointly transmitted via the IntL output pin. Relevant advertisement, threshold, and readout registers are found in the SFF-8636 MSA.

4 Specifications

4.1 Absolute Maximum Specifications

Absolute maximum ratings are those beyond which damage to the device may occur.

Prolonged operation between the operational specifications and absolute maximum ratings is not intended and may cause permanent device degradation.

Parameter	Symbol	Min	Max	Units
Storage Temperature	T_S	-40	85	°C
Operating Case Temperature	T_{OP}	0	70	°C
Supply Voltage	Vcc	-0.5	3.6	V
Relative Humidity (non-condensing)	RH – Option 1	5	95	%
Control Input Voltage	V_I	-0.3	Vcc+0.5	V



Module temperature per DDMI readout of up to 75°C is allowed.

4.2 Recommended Operating Conditions and Power Supply Requirements

Parameter	**Symbol**	**Min**	**Typ**	**Max**	**Units**
Power Supply Voltage	VCC	3.135	3.3	3.465	V
Instantaneous peak current at hot plug (400G)	ICC_IP	-	-	3600	mA
Sustained peak current at hot plug (400G)	ICC_SP	-	-	3000	mA
Maximum Power consumption (400G)	PD	-	8.1	8.5	W
Maximum Power consumption, Low Power Mode (400G)	PDLP	-	-	2	W
Instantaneous peak current at hot plug (200G)	ICC_IP	-	-	2200	mA
Sustained peak current at hot plug (200G)	ICC_SP	-	-	1840	mA
Maximum Power consumption (200G)	PD	-	-	5.5	W
Maximum Power consumption, Low Power Mode (200G)	PDLP	-	-	2	W
Signaling Rate per Lane	SRL	-	53.125	-	GBd
Two Wire Serial Interface Clock Rate	-	-	-	400	kHz
Power Supply Noise Tolerance (10Hz - 10MHz)	-	66	-		mV
Rx Differential Data Output Load	-	-	100	-	Ohm
Operating distance (OM3)		2		30	m
Operating distance (OM4)		2		50	m

4.3 Electrical Specifications

Parameter	Min	Typ	Max	Units
Receiver (Module Input)				
AC common-mode output Voltage (RMS)	-	-	25	mV
Differential output Voltage (Long mode)	-	-	845	mV
Differential output Voltage (Short mode)	-	-	600	mV
Near-end Eye height, differential	70	-	-	mV
Far-end Eye height, differential	30	-	-	mV
Far end pre-cursor ratio	-4.5	-	2.5	%
Differential Termination Mismatch	-	-	10	%
Transition Time (min, 20% to 80%)	9.5	-	-	ps
DC common mode Voltage	-350	-	2850	mV
Transmitter (Module Input)				
Differential pk-pk input Voltage tolerance	750	-	-	mV
Differential termination mismatch	-	-	10	%
Single-ended voltage tolerance range	-0.4	-	3.3	V
DC common mode Voltage	-350	-	2850	mV

Notes:

Amplitude customization beyond these specs is dependent on validation in customer system.

4.3.1 Electrical Specification for Low Speed Signal

Parameter	Symbol	Min	Max	Units
Module output SCL and SDA	VOL	0	0.4	V
	VOH	VCC-0.5	VCC+0.3	V
Module Input SCL and SDA	VIL	-0.3	VCC*0.3	V
	VIH	VCC*0.7	VCC+0.5	V

4.4 Optical Specifications

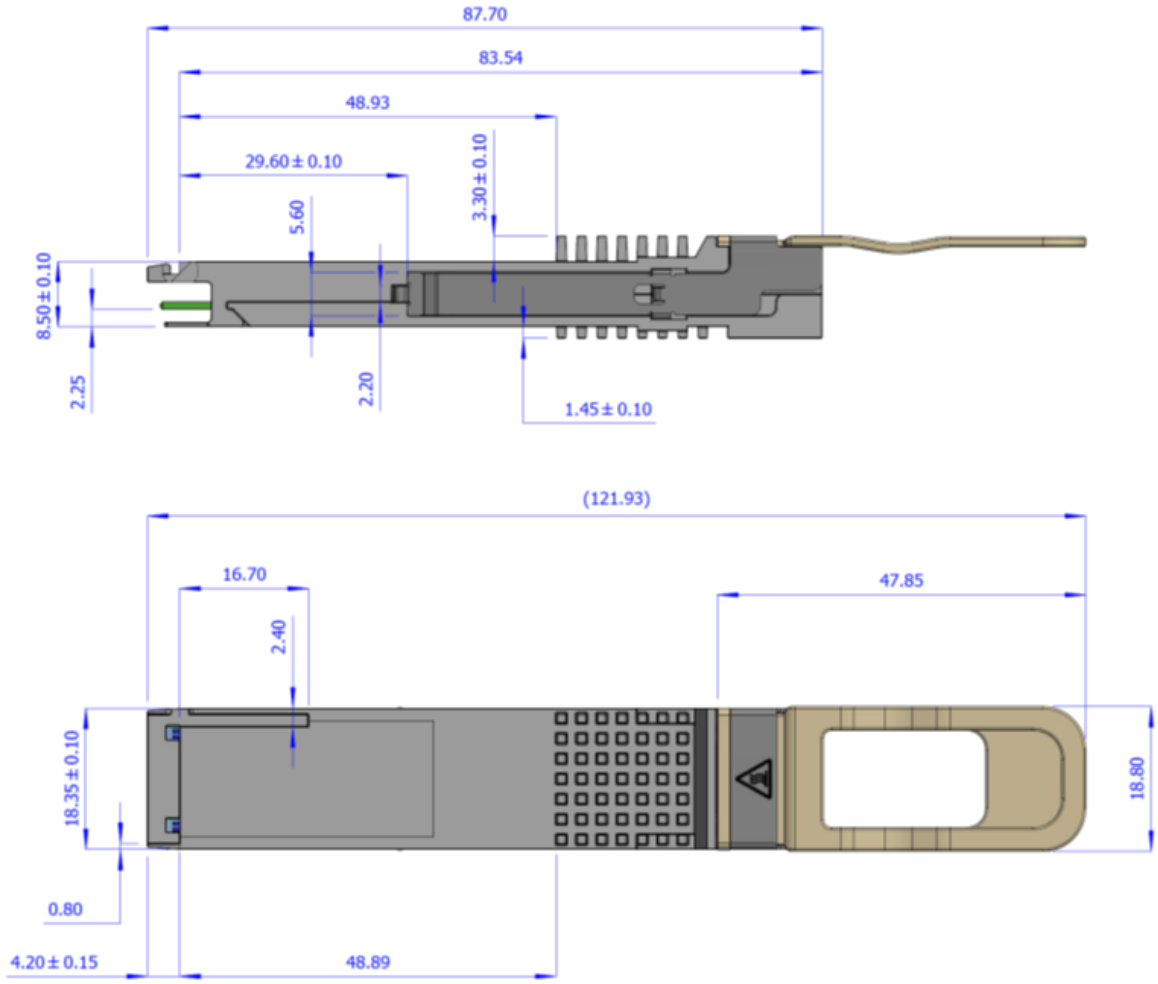
Parameter	Symbol	Min	Typ	Max	Unit	Notes
Transmitter						
Wavelength	λC	844	850	863	nm	
RMS spectral width	Δλ			0.6		
Average Launch Power, each lane	AOPL	-4.6	-	4.0	dBm	1

Outer Optical Modulation Amplitude (OMAouter), each lane (min)	TOMA	-2.6		3.5	dBm	2
Transmitter and Dispersion Eye Closure for PAM4 (TDECQ), each lane	TDECQ	-	-	4.4	dB	
Average Launch Power of OFF Transmitter, each lane	TOFF	-	-	-30	dBm	
Extinction Ratio, each lane	ER	-	2.5		dB	
RIN21.4OMA	RIN	-	-	-132	dB/ Hz	
Optical Return Loss Tolerance	ORL	-	-	12	dB	
Transmitter Reflectance	TR	-	-	-26	dB	3
Receiver						
Wavelength	λC	842	850	863	nm	
Damage Threshold, average optical power, each lane	AOPD	5	-	-	dBm	
Average Receive Power, each lane	AOPR	-6.3	-	4.0	dBm	6
Receive Power (OMAouter), each lane	OMA-R	-	-	3.5	dBm	
Receiver Reflectance	RR	-	-	-26	dB	
Receiver Sensitivity (OMAouter), each lane	SOMA	-	-	-4.4	dBm	4
Stressed Receiver Sensitivity (OMAouter), each lane	SRS	-	-	-1.8	dBm	5
Conditions of stressed receiver sensitivity test						
Stressed eye closure for PAM4	SECQ	4.4		dB		
OMAouter of each aggressor lane	OMAouter	3.5		dBm		

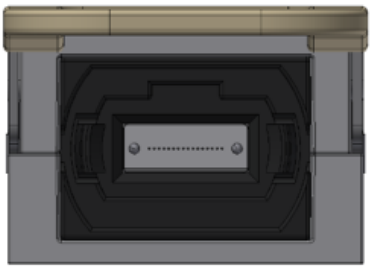
Notes:

1. Average launch power, each lane (min) is informative and not the principal indicator of signal strength.
2. Even if max(TECQ,TDECQ) < 1.8dB, OMAouter (min) must exceed this value.
3. Transmitter reflectance is defined looking into the transmitter.
4. Receiver sensitivity (OMAouter), each lane (max) is informative and is defined for a transmitter with SECQ of 0.9 dB.
5. Measured with conformance test signal at TP3 for the BER = 2.4x10-4
6. Minimum power is informative. AOP above the minimum does not ensure compliance

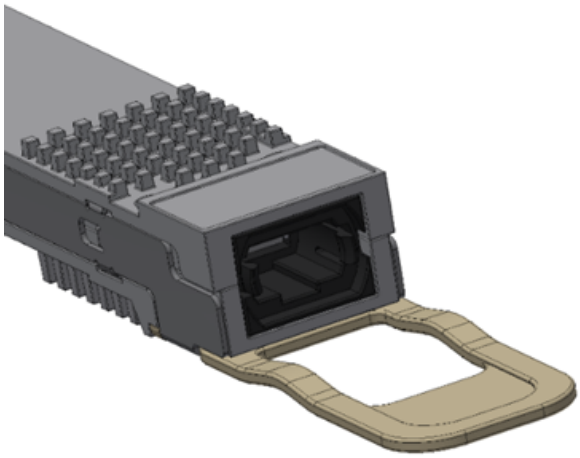
Mechanical Specifications



Connector illustration



This is the ‘head end’. The connectors in the fiber cable splitter ‘tails’ have same layout but only 2 transmit and 2 receive fibers with the middle 8 positions empty.



Module port labeling and lane routing. Txn/Rxn refers to the OSFP pin description

Single mode optics are denoted by a yellow-colored pull tab and yellow-colored optical fiber. The green plastic shell on the MPO-12/APC optical connector denotes Angled Polish Connector.

> ⚠ Images are for illustration purposes only. Product labels, colors, and form may vary.

4.4.1 New Form-factors

The QSFP112 body has a flat-top to cool the 8-Watt NDR transceiver. Unlike the Twin-port transceiver, the heat sink is contained on the ConnectX-7 or BlueField-3 card connector cage. The small bumps on top of the transceiver, accelerate cooling but remain outside the QSFP112 cage.

4.5 Labels

4.5.1 Back shell Label

The label applied on the transceiver's back-shell is illustrated below. Note that the Images are for illustration purposes only. Labels look and placement may vary.

Transceiver Label (Illustration)



Images are for illustration purposes only. Product labels, colors, and form may vary.

4.5.1.1 Transceiver Back-Shell Label Serial Number Legend

Symbol	Meaning	Notes
MT	Manufacturer name	2 digits (alphanumeric)
YY	Year of manufacturing	2 last digits of the year (numeric)
WW	Week of manufacturing	2 digits (numeric)
JC *or* DM	Manufacturer Site: JC – Option 1 (China) DM – Option 2 (Malaysia)	Two characters
SSSSS	Serial number	5 digits (decimal numeric) for serial number, starting from 00001.

4.5.2 Regulatory Compliance

The transceiver is a Class 1M laser product. It is certified per the following standards:

Feature	Agency	Standard
Laser Eye Safety	FDA/CDRH	CDRH 21 CFR 1040 and Laser Notice 50
Electrical Safety	CB	IEC 62368
Electrical Safety	UL/CSA	UL 62368 and CAN/CSAN 62368



Warning: Exposure to the laser light can cause damage to the eyes. Use protective face gear while handling this product and keep the laser light away from the eyes and face.

4.6 Connector and Cabling Details

4.6.1 MPO-12/APC Optical Connector

The Twin-port NDR transceiver has a unique NVIDIA patented design enabling two, multiple-push-on/angled-polished-connector 12-fiber (MPO-12/APC) optical connectors per single OSFP form-factor by turning the optical connectors vertically in the twin-port transceiver end. This enables it to host two NDR transceivers inside, each with its own MPO-12/APC optical connector operating independently that can link to another Twin-port transceiver or to a single-port 400Gb/s NDR transceiver.

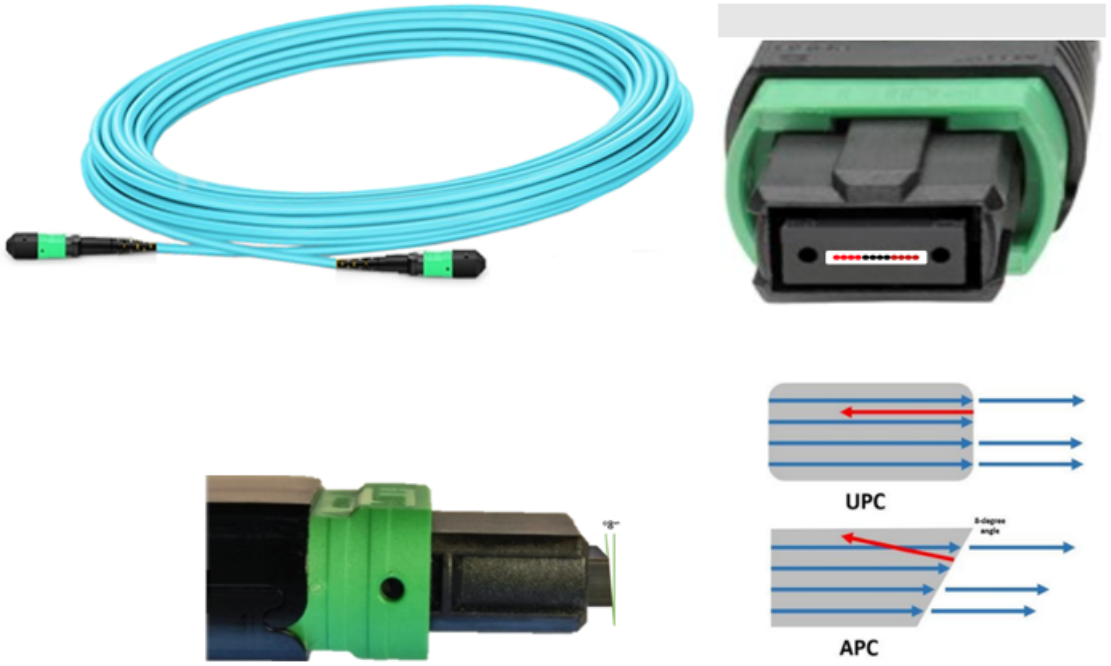
The MPO-12 has a 12-fiber ribbon but only 8-fibers are used – four transmit and four receive fibers for the 4-channels of 100G-PAM4 NDR.

- The APC design minimizes back reflections and signal interference by diverting back reflected light from the fiber face to be absorbed into the fiber cladding.
- A positioning key on top of the connector together with the alignment pins define the fiber position numbering scheme to align pin 1 in the optical connector to pin 1 in the transceiver also called “polarity”
- Transceivers have alignment pins for precise positioning of the cable connector against the optical beams. The fiber cable has alignment holes matching the transceiver’s pins.
- It is important to note that transceivers have pins. Optical connectors have holes and are used with transceivers. Optical connectors with pins are not compatible with transceivers and are used in trunk cabling to connect two fiber cables together.

The MPO-12/APC optical connector is used in both the NDR single mode and multimode fiber cables.

Multimode optics is denoted by a tan-colored pull tab and aqua-colored optical fiber. Green plastic shell on the MPO-12/APC connector denotes Angled Polish Connector and is not compatible with aqua colored shell for Ultra-flat Polished Connectors (UPC) for HDR.

MPO-12/APC Showing 4-Transmit and 4-Receive Fibers and Angled Polish Connector End



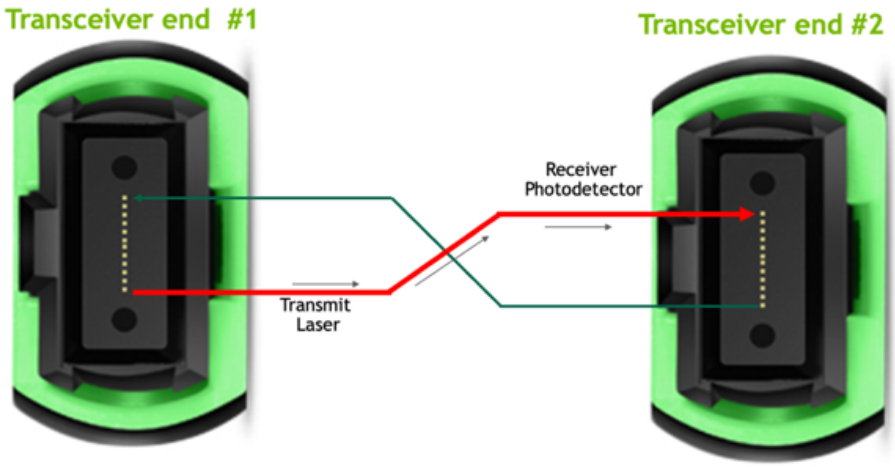
4.6.1.1 NVIDIA Supplied Crossover Type-B Fiber Cables

Linking two transceivers directly together requires aligning the transceiver laser sources with the correct photo detectors in the receive transceiver. Transmit and receive fibers are switched inside the cable enabling two transceivers to be directly connected to each other. This is called a Type-B crossover fiber.

Each of the two 4-channel NDR ports in the Twin-port transceiver has its own 4-channel optical connector that can link to two single-port 400Gb/s NDR transceiver. Two fiber cables are needed for each Twin-port transceiver.

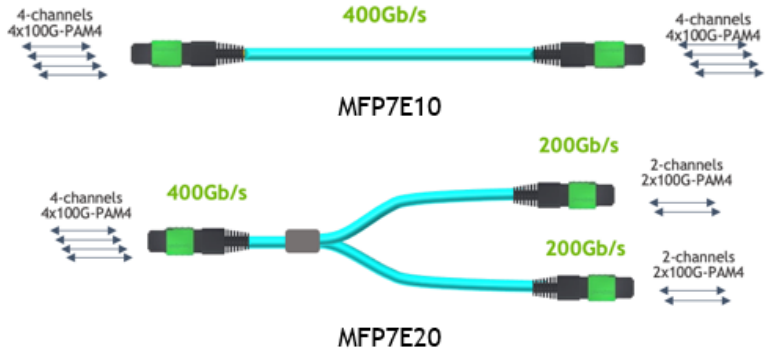
Fiber cables are crossover cable Type-B that aligns the transmit laser with the opposite transceiver's receiver photodetector allowing to directly connect two transceivers together to maintain minimum optical losses, lowest back reflections, longest reach and increased reliability without the use of optical patch panels. For Twin-port transceivers, both cables must be the same type (straight or 1:2 splitter) although different lengths are allowed.

Crossover Cable Design



Multimode Straight and Splitter Cables

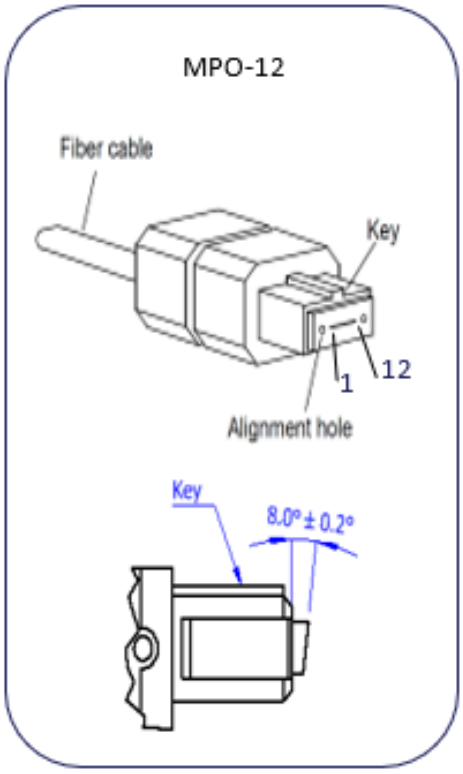
Twin-port transceiver side NDR or NDR200 HCA DPU side

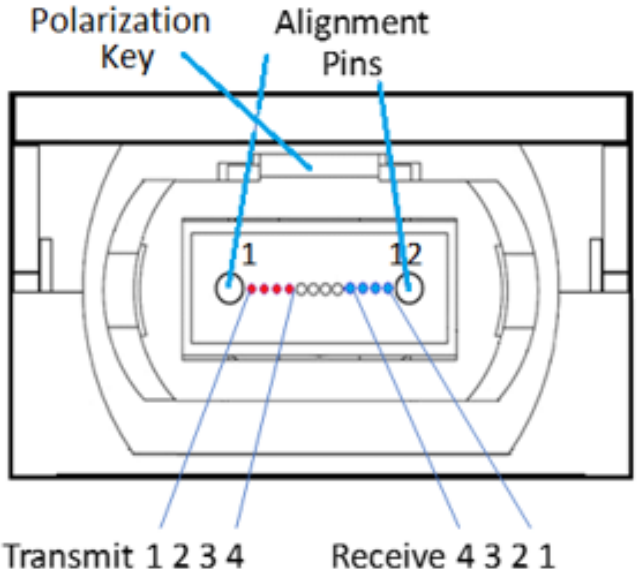


Note: Refer to the Recommended Fiber Cables table for more information.

Transceivers have alignment pins for precise positioning of the cable connector against the optical beams. The fiber cable has alignment holes matching the transceiver's pins.

MPO Connector with Alignment Holes and Positioning Key





Reference: IEC specification IEC 61754-7

NDR transceiver: MPO Receptacle, Lane Assignment, and Positioning Key (front view)

4.6.2 Handling and Cleaning

The transceiver can be damaged by exposure to current surges and over voltage events. Take care to restrict exposure to the conditions defined in Absolute Maximum Ratings. Observe normal handling precautions for electrostatic discharge-sensitive devices.

The transceiver is shipped with dust caps on both the electrical and the optical port. The cap on the optical port should always be in place when there is no fiber cable connected. The optical connector has a recessed connector surface which is exposed whenever it has no cable nor cap.

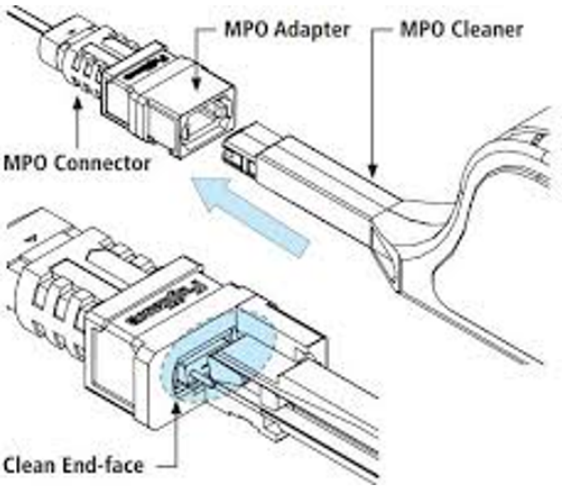
Important note 1: *Keep both the fiber and transceiver dust caps.*

Important note 2: *Clean both transceiver receptacle and cable connector* prior to insertion of the fiber cable, to prevent contamination from it.

The dust cap ensures that the optics remain clean during transportation. Standard cleaning tools and methods should be used during installation and service. Liquids must not be applied.

Important note 3: *80% of transceiver link problems are related to dirty optical connectors*.





4.6.3 Cable Management Guidelines

For more information and general interconnect management and installation, see NVIDIA Cable Management Guidelines and FAQ Application Note.

5 Ordering Information

5.1 Part Numbers and Description

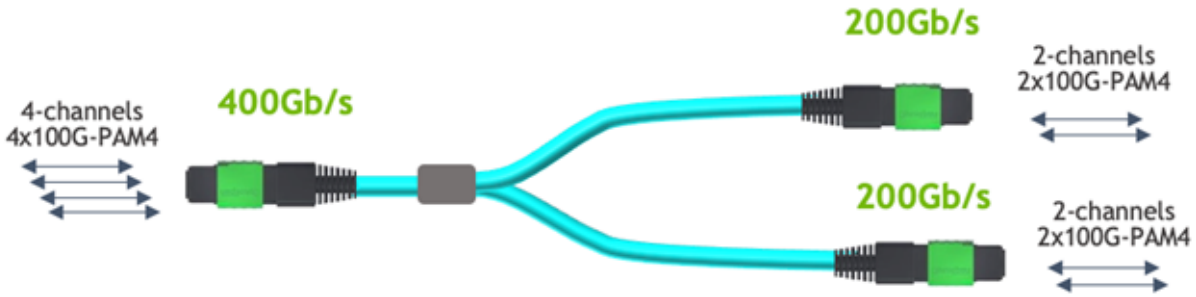
Legacy OPN	NVIDIA OPN	Description
MMA1Z00-NS400	980-9I693-00NS00	NVIDIA single port transceiver, 400Gbps,NDR, QSFP112, MPO12 APC, 850nm MMF, up to 50m, flat top

6 Recommended NVIDIA Supplied Crossover Fiber Cables Part Numbers



Multimode, Straight Crossover Fibers

OPN	4-channel MPO/APC to 4-channel MPO/APC
MFP7E10-N003	3m
MFP7E10-N005	5m
MFP7E10-N007	7m
MFP7E10-N010	10m
MFP7E10-N015	15m
MFP7E10-N020	20m
MFP7E10-N030	30m
MFP7E10-N050	50m



Multimode, 1:2 Splitter Crossover Fibers

OPN	4-channel MPO/APC to Two 2-channel MPO/APC
MFP7E20-N003	3m
MFP7E20-N005	5m
MFP7E20-N007	7m
MFP7E20-N010	10m
MFP7E20-N015	15m

OPN	4-channel MPO/APC to Two 2-channel MPO/APC
MFP7E20-N020	20m
MFP7E20-N030	30m
MFP7E20-N050	50m

7 Document Revision History

Rev	Date	Description
1.7	May 2024	Modified max power for 200G (2 channels) to 5.5W (was 6.5W)
1.6	Mar. 2024	Added DDMI note in the Specifications section.
1.5	Jan. 2024	Updated low power mode output.
1.4	Dec. 2023	Updated Regulatory Compliance section.
1.3	Aug. 2023	Updated maximum power consumption.
1.2	Apr. 2023	▪ Updated the document for Ethernet support. ▪ Minor text edits
1.1	Feb. 2023	Updated Regulatory Compliance section.
1.0	Nov. 2022	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 its affiliates 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
© 2026 NVIDIA Corporation & affiliates. All Rights Reserved.

NVIDIA