



**MMA4Z00-NS 800Gb/s Twin-port OSFP, 2x400Gb/s
Multimode 2xSR4, 50m**

Table of contents

Overview	5
Pin Description	11
Specifications	16
Part Numbers and Description	32
Recommended NVIDIA Supplied Crossover Fiber Cables Part Numbers	33
Document Revision History	35

Introduction

The NVIDIA MMA4Z00-NS is an InfiniBand and Ethernet 800Gb/s 2x400Gb/s Twin-port OSFP, SR8 (2xSR4) multimode, parallel, 8-channel transceiver using two, 4-channel MPO-12/APC optical connectors at 400Gb/s each. The parallel multimode, short reach 8-channel (2xSR4) uses 100G-PAM4 modulation and has a maximum fiber reach of 50-meters using 8 multimode fibers. The 50-meter length assumes two optical patch panels in the link.


The Twin-port 2xSR4 transceiver is a key NVIDIA innovation with two internal transceiver engines enabling 64-ports of 400Gb/s in a 32-OSFP cage Quantum-2 switch. Spectrum-4 switches have 32 or 64 cages and enable 64-128 400G ports. The transceiver firmware supports both InfiniBand and Ethernet and is automatically enabled based on the switch protocol.

The Quantum-2 and Spectrum-4 switches require finned-top OSFP shells for extra transceiver cooling. The main application for MMA4Z00-NS is linking two switches together with up to 50-meter transceivers (marked blue).

A flat-top version offered for liquid-cooled and DGX H100 Cedar7 systems links.

The transceiver combinations guarantee optimal operation in NVIDIA end-to-end systems and customer networking solutions. Rigorous production testing ensures the best out-of-the-box installation experience, performance, and durability.

Finned Top Transceiver



Transceiver

Flat Top



i Note

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

Key Features

- InfiniBand and Ethernet
- 800G 2xSR4 multimode transceiver
- 8-channels of 100G-PAM4 electrical modulation
- Two ports of 4-channel 100G-PAM4 optical modulation
- Supports two straight 400Gb/s or two 1:2 splitter fiber cables for 200Gb/s
- Finned-top OSFP for air-cooled switches
- Flat-top OSFP for DGX- H100 or liquid-cooled

- 850nm VCSEL
- Maximum reach:
 - 30m using OM3 fiber
 - 50m using OM4 fiber
- Two MPO-12/APC optical connectors
- 17 Watts max power
- 2.0 Watt low-power sleep mode
- Single 3.3V power supply
- Class 1 laser safety
- Hot pluggable, RoHS compliant
- [OSFPmsa.org](https://www.osfpmsa.org) compliant
- CMIS 4.0 compliant
- Case temperature range 0°C to +70°C

Applications

- Used in Quantum-2 air-cooled switches, liquid-cooled and DGX-H100 systems to switches, ConnectX-7, and BlueField-3 DPUs

Overview

The twin-port transceiver is a key innovation for expanding the 400G NDR InfiniBand Quantum-2 and 400GbE Spectrum-4 Ethernet (400G IB/EN) switch capabilities. The line rate is 400Gb/s for both 400GbE Ethernet and NDR InfiniBand based on the 100G-PAM4 modulation. The switches use Twin port OSFP cages supporting two transceiver engines in a single OSFP form-factor plug creating 800Gb/s electrical to the switch and 2x400G optics using two MPO-12/APC optical connectors. Both switches use the same Twin port, 2x400G OSFP plugs for transceivers, copper DACs and ACCs, and are only used in Quantum-2 and Spectrum-4 OSFP air-cooled switches.

Twin-port devices enable several unique configurations to connect switches, adapters and DPUs. Switches maintain separate protocols. The MMA4Z00-NS transceiver includes both InfiniBand and Ethernet protocols which is activated depending on the switch protocol the transceiver used.

Twin-port Transceiver Connectivity Scenarios

The twin-ports enable several unique configurations to connect switches to switches, ConnectX-7 adapters, and BlueField-3 DPUs.

The Twin port OSFP uses two, 4-channel MPO-12/APC optical connectors with two 4-channel fiber cables. These can link to a single port 400G OSFP or QSFP112 transceivers used in ConnectX-7 adapters and/or BlueField-3 DPUs.

The electronics, optics and optical connectors are the same for both single port OSFP and QSFP112. Both ConnectX-7 and BlueField-3 devices can be used with Twin port OSFP transceivers at the same time with their respective form-factor type.

- Twin-port multimode OSFP transceivers remain at 15 Watts for all configurations linking OSFP switches, OSFP and QSFP112 adapters, and DPUs simultaneously.
- Linking Twin-port transceivers with 1:2 fiber splitter cables to 400G transceivers automatically creates 200Gb/s transceivers by activating only 2-channels and automatically reduces power consumption in the 400G transceiver from 8 to 5.5 Watts.
- Twin-Port transceivers require ordering two fibers at specific lengths required

- Each fiber cable can each be different lengths. Both fibers need to be same type: both straight or both splitters, and not mixed.
- Both fibers should be the approximately same length to avoid inducing different latency delays in the fibers (4.5ns/meter).
- Straight and splitter fibers cannot be used at the same time in Twin port OSFP transceivers.
- NVIDIA supplies multimode, crossover, straight fiber cables up to 100-meters straight and 50-meters for splitters that enable linking transceivers directly together.

All combinations of Twin port 2x400G OSFP, 400G single port OSFP/QSFP112, ConnectX-7 and BlueField-3 contains both InfiniBand and Ethernet protocols which is activated upon inserting into an InfiniBand or Ethernet switch. This enables one set of cables, transceivers, adapters and DPUs to have multiple uses in a network – especially in DGX systems where low-latency InfiniBand switching may be used for GPU-to-GPU networking and Ethernet switching systems for storage and cluster communication.

The use cases include:

1. Switch-to-switch at 800Gb/s or to two switches at 400Gb/s
2. Switch-to-two combinations of 400G ConnectX-7 BlueField-3
3. Switch-to-four combinations of 200G ConnectX-7 and/or BlueField-3
4. Switch linked to the DGX H100 “Viking” CPU chassis to Cedar-7 complex

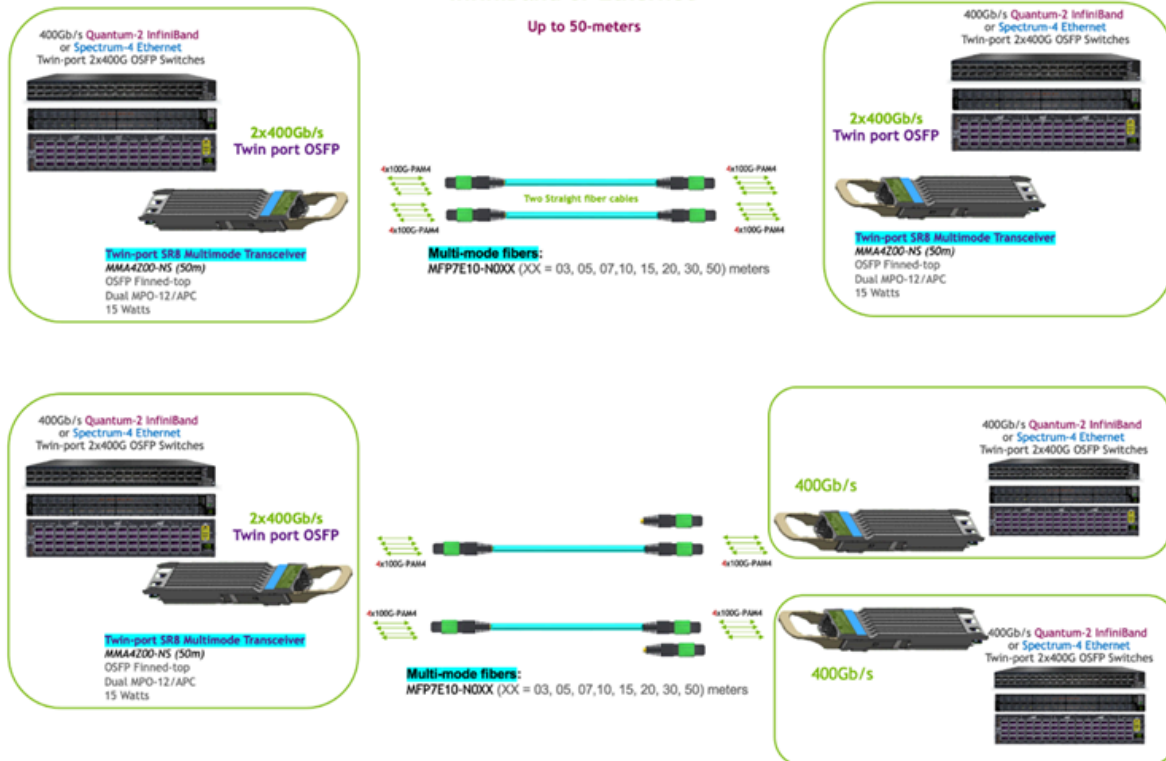
1. **Switch-to-switch at 800Gb/s or to two switches at 400Gb/s**

Two OSFP-based switches can be linked together at 800G (2x400G) using two Twin port OSFP transceivers and two straight, multimode fiber cables (MFP7E10-Nxxx) up to 50-meters. Also, the two fiber cables can be routed to two different switches forming two 400Gb/s links. The extra Twin port OSFP ports can be then routed to more switches.

400G IB/EN SWITCH-TO-SWITCH OSFP LINKS

Multimode: 2x400G Twin-Port OSFP-to- 2x400G Twin port OSFP

InfiniBand or Ethernet



2. Switch-to-two combinations of 400G ConnectX-7 and/or BlueField-3

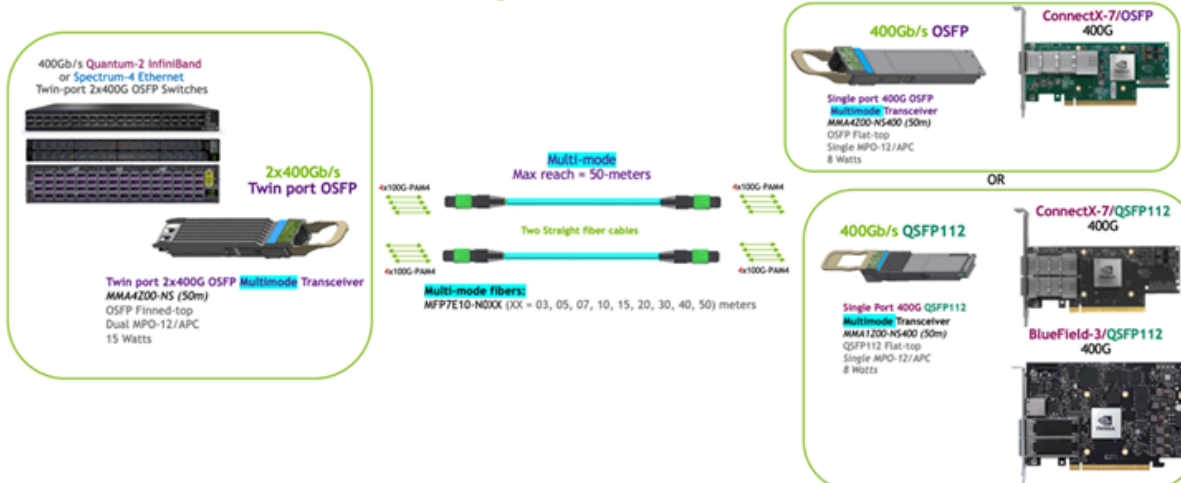
A Twin port OSFP transceiver using two, straight fiber cables can support up to two adapter and/or DPU combinations. Each of the two, 4-channel fiber cables (MFP7E10-N0xx) can link to a 400G transceiver up to 50-meters in either OSFP (MMA4Z00-NS400) or QSFP112 (MMA1Z00-NS400) form-factor. The electronics, optics and optical connectors are the same for both single port OSFP and QSFP112 form-factors and power consumption is 8 Watts.

- Only ConnectX-7/OSFPs supports single port OSFP.
- QSFP112 form-factor is used in ConnectX-7/QSFP112s and/or BlueField-3/QSFP112 DPUs.

Any combination of ConnectX-7 and BlueField-3 using OSFP or QSFP112 can be used at the same time on Twin port OSFP transceivers.

400G IB/EN SWITCH-TO- 2 CONNECTX-7 AND BLUEFIELD-3

Multimode: 2x400G Twin-Port -to- ConnectX-7/OSFP, ConnectX-7/QSFP112 or BlueField-3/QSFP112. InfiniBand or Ethernet



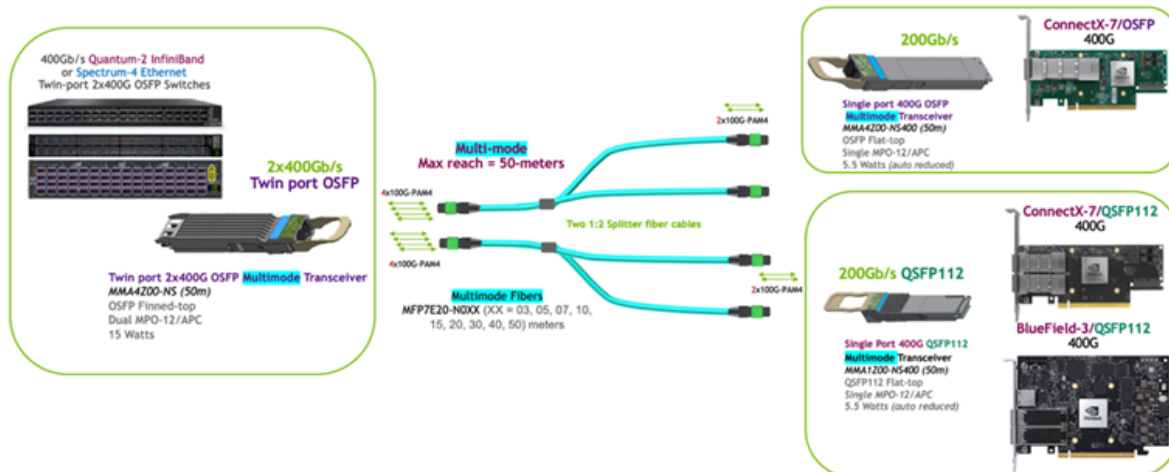
3. Switch-to-four combinations of 200G ConnectX-7 and/or BlueField-3

A Twin port OSFP transceiver using two, 1:2 fiber splitter cables can support up to four adapter and/or DPU combinations. Each of the two, 4-channel 1:2 fiber splitter cables (MFP7E20-N0xx) can link to a 400G transceiver up to 50-meters in either OSFP (MMA4Z00-NS400) or QSFP112 (MMA1Z00-NS400) form-factor. The electronics, optics and optical connectors are the same for both single port OSFP and QSFP112 form-factors. 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 to 5.5 Watts. Twin port OSFP power consumption remains at 15 Watts.

- Only ConnectX-7/OSFPs can accept single port OSFP.
- QSFP112 form-factor is used in ConnectX-7/QSFP112s and/or BlueField-3/QSFP112 DPUs.
- Any combination of ConnectX-7 types and BlueField-3 can be used at the same time on Twin port OSFP transceivers.

400G IB/EN SWITCH-TO- 4 CONNECTX-7 AND BLUEFIELD-3

Multimode: 2x400G Twin-Port -to- 200G ConnectX-7/OSFP, ConnectX-7/QSFP112 or BlueField-3/QSFP112. InfiniBand and/or Ethernet



4. Switch-to-DGX H100 GPU Systems

The DGX-H100 contains eight “Hopper” H100 GPUs in the top chassis section, two CPUs, storage, and InfiniBand and/or Ethernet networking in the bottom server section. This contains eight 400Gb/s ConnectX-7 ICs mounted on two mezzanine boards called “Cedar-7” cards for GPU-to-GPU InfiniBand or Ethernet networking. The cards I/Os are routed internally to four 800G Twin-port OSFP cages with internal riding heat sinks on top of the cages mounted on the front panel. This requires the use of flat-top transceivers, ACCs, and DACs in the DGX H100. The 400G IB/EN switches require finned-top 2x400G transceivers for additional cooling due to the reduced air flow inlets in the switches.

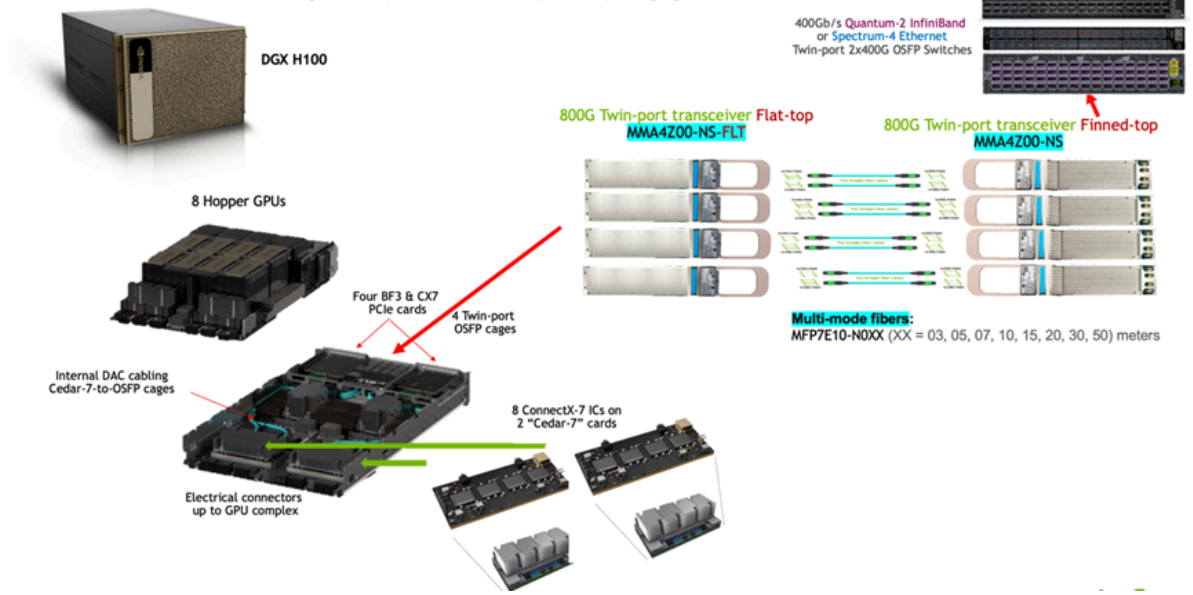
The Cedar-7-to-Switch links can be either single mode or multimode optics or ACC active copper cables and in InfiniBand or Ethernet.

Each Twin-port 2x400G transceiver provides two 400G ConnectX-7 links from the DGX to the Quantum-2 or Spectrum-4 switch. This reduces the ConnectX-7 card redundancy, complexity, and the number of transceivers compared to the DGX A100, which uses 8 separate HCAs and 8 transceivers or AOCs and two additional ConnectX-6s for InfiniBand or Ethernet Storage.

Additionally, for traditional networking to storage, clusters, and management, the DGX-H100 also supports up to four ConnectX-7 and/or two BlueField-3 DPUs in InfiniBand and/or Ethernet for storage I/O, and additional networking using 400G or 200G with OSFP or QSFP112 devices. These PCIe card slots are located on both sides of the OSFP GPU cages and use separate cables and/or transceivers.

400G IB/EN SWITCH-TO-DGX H100/CEDAR-7 LINKS

Using Four Twin-port 800G Transceivers (Multimode) linking Eight ConnectX-7 ICs @ 400G



Pin Description


The device is OSFP MSA Specification for OSFP Octal Small Form Factor Pluggable Module Rev. 1.12 compliant, see www.osfpmsa.org.

OSFP Pin Description

Pin	Symbol	Description	Pin	Symbol	Description
1	GND	Ground	31	GND	Ground
2	Tx2p	Transmitter Non-Inverted Data Input	32	Rx2p	Receiver Non-Inverted Data Output
3	Tx2n	Transmitter Inverted Data Input	33	Rx2n	Receiver Inverted Data Output
4	GND	Ground	34	GND	Grounds
5	Tx4p	Transmitter Non-Inverted Data Input	35	Rx4p	Receiver Non-Inverted Data Output
6	Tx4n	Transmitter Inverted Data Input	36	Rx4n	Receiver Inverted Data Output
7	GND	Ground	37	GND	Ground
8	Tx6p	Transmitter Non-Inverted Data Input	38	Rx6p	Receiver Non-Inverted Data Output
9	Tx6n	Transmitter Inverted Data Input	39	Rx6n	Receiver Inverted Data Output
10	GND	Ground	40	GND	Ground
11	Tx8p	Transmitter Non-Inverted Data input	41	Rx8p	Receiver Non-Inverted Data Output
12	Tx8n	Transmitter Inverted Data Input	42	Rx8n	Receiver Inverted Data Output
13	GND	Ground	43	GND	Ground
14	SCL	2-wire serial interface clock	44	INT / RSTn	Module Interrupt / Module Reset

Pi n	Symbol	Description	Pi n	Symbol	Description
15	VCC	+3.3V Power	45	VCC	+3.3V Power
16	VCC	+3.3V Power	46	VCC	+3.3V Power
17	LPWn / PRSn	Low-Power Mode / Module Present	47	SDA	2-wire Serial interface data
18	GND	Ground	48	GND	Ground
19	Rx7n	Receiver Inverted Data Output	49	Tx7n	Transmitter Inverted Data Input
20	Rx7p	Receiver Non-Inverted Data Output	50	Tx7p	Transmitter Non-Inverted Data Input
21	GND	Ground	51	GND	Ground
22	Rx5n	Receiver Inverted Data Output	52	Tx5n	Transmitter Inverted Data Input
23	Rx5p	Receiver Non-Inverted Data Output	53	Tx5p	Transmitter Non-Inverted Data Input
24	GND	Ground	54	GND	Ground
25	Rx3n	Receiver Inverted Data Output	55	Tx3n	Transmitter Inverted Data Input
26	Rx3p	Receiver Non-Inverted Data Output	56	Tx3p	Transmitter Non-Inverted Data Input
27	GND	Ground	57	GND	Ground
28	Rx1n	Receiver Inverted Data Output	58	Tx1n	Transmitter Inverted Data Input
29	Rx1p	Receiver Non-Inverted Data Output	59	Tx1p	Transmitter Non-Inverted Data Input
30	GND	Ground	60	GND	Ground

OSFP Module Pad Layout



The Active Optical Cable (AOC) pin assignment is SFF-8679 compliant.

Control Signals (OSFP)

This device supports CMIS 4.0 (check for update, e.g. to CMIS 5) compliant management interface and OSFP MSA compliant form factor and interfaces. This implies that the control signals shown in the pad layout are implemented with the following functions:

Name	Function	Description
LPWn/PRSn	Input/output	Multi-level signal for low power control from host to module and module presence indication from module to host. This signal requires the circuit as described in the OSFP Specification [].
INT/RSTn	Input, /outp	Multi-level signal for interrupt request from module to host and reset control from host to module. This signal requires the circuit as described

Name	Function	Description
	ut	in the OSFP Specification [].
SCL	BiDir	2-wire serial clock signal. Requires pull-up resistor to 3.3V on host.
SDA	Bidir	2-wire serial data signal. Requires pull-up resistor to 3.3V on host.

Diagnostics and Other Features

The transceiver has a microcontroller with functions for monitoring supply voltage, temperature, laser bias current, optical transmit and receive levels with associated warning and alarm thresholds that can be read by the switch software and viewed remotely.

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

Physical layer link optimization:

- Adaptive Tx input equalization
- Programmable Rx output amplitude
- Programmable Rx output pre-cursor
- Programmable Rx output post-cursor

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 programmable)

Page 13h and 14h Module Diagnostics

- Host side and line side loopback

- PRBS generator and checker on host and line interfaces

Interrupt indications:

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

Other CMIS 4.0 functions

FW upgrade supported via CDB commands.

Specifications

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.

Note

- For Ethernet end-to-end, the transceiver performance is guaranteed above 15°C.
- Module temperature per DDMI readout of up to 75°C is allowed.
- Maximum switch ambient temperature for forward (back to front) airflow on QM9700 is 35°C with all fans, and 30°C in case of fan failure.

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	ICC_IP	-	-	6800	mA
Sustained peak current at hot plug	ICC_SP	-	-	5670	mA
Maximum Power Dissipation	PD	-	15	17	W
Maximum Power Dissipation, Low Power Mode	PDLP	-	-	2	W
Signaling Rate per Lane	SRL	-	53.125	-	GBd
Two Wire Serial Interface Clock Rate	-	100	-	1000	kHz

Parameter	Symbol	Min	Typ	Max	Units
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

Electrical Specifications

Parameter	Symbol	Min	Typ	Max	Units
Receiver (Module Output)					
AC common-mode output Voltage (RMS)		-	-	25	mV
Differential output Voltage (Long mode)		-	-	845	mV
Differential output Voltage (Short mode)		-	-	600	mV
Eye height, differential		15	-	-	mV
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.

Electrical Specification for Low Speed Signal

Parameter	Symbol	Min	Max	Units
Module output SCL and SDA	VOL	0	0.4	V

Parameter	Symbol	Min	Max	Units
	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

Optical Specifications

Parameter	Symbol	Min	Typ	Max	Unit	Notes
Transmitter						
Wavelength	λ_C	844	850	863	nm	
RMS spectral width	DI			0.6		
Average Launch Power, each lane	AOPL	-1.0	-	4.0	dBm	1
Outer Optical Modulation Amplitude (OMA _{outer}), each lane (min)	TOMA	-2.1		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		3.5	dB	
RIN14OMA	RIN	-	-148		dB/Hz	
Optical Return Loss Tolerance	ORLT		-	14	dB	
Transmitter Reflectance	TR	-	-	-26	dB	3
Receiver						
Wavelength	λ_C	842	850	863	nm	
Damage Threshold, average optical power, each lane	AOPD	5	-	-	dBm	
Average Receive Power, each lane	AOPR	-6.	-	4.0	dBm	6

Parameter	Symbol	Min	Typ	Max	Unit	Notes
		3				
Receive Power (OMA _{outer}), each lane	OMA-R	-	-	3.5	dBm	
Receiver Reflectance	RR	-	-	-20	dB	
Receiver Sensitivity (OMA _{outer}), each lane	SOMA	-	-	-4.6	dBm	4
Stressed Receiver Sensitivity (OMA _{outer}), each lane	SRS	-	-	-2.0	dBm	5
Conditions of stressed receiver sensitivity test						
Stressed eye closure for PAM4	SECQ	4.4			dB	
OMA _{outer} of each aggressor lane	OMA _{outer}	3.5			dBm	

Notes:


1. Average launch power, each lane (min) is informative and not the principal indicator of signal strength.
2. Even if $\max(\text{TECQ}, \text{TDECQ}) < 1.8\text{dB}$, OMA_{outer} (min) must exceed this value.
3. Transmitter reflectance is defined looking into the transmitter.
4. Receiver sensitivity (OMA_{outer}), each lane (max) is informative and is defined for a transmitter with $\text{TDECQ} \leq 1.8\text{ dB}$
5. Measured with conformance test signal at TP3 for the $\text{BER} = 2.4 \times 10^{-4}$

Minimum power is informative. AOP above the minimum does not ensure compliance

Mechanical Specifications


Finned-top twin-port for air-cooled 400G IB/EN Switches: Bottom, Side, and Top Views.

Option 1:




Enlarged view of detail A and B


Transceiver Port View



Option 2:




Option 3:



i Note

1. Front of the module. Pull tab and other components can extend 1.6mm max from the bottom of the module and can have up to 9.3mm width in the max envelope shown.
2. Applies to the type 1 module only.
3. Indicated surfaces (all 4 sides) to be conductive for connection to chassis ground.
4. Applies from the top of the module to the bottom of the module, Inside the cage.

**Flat-top Twin-port transceiver for DGX-H100 and Liquid-cooled Systems:
Bottom, Side, and Top Views.**




Note

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

Labels

Transceiver Labeling and Fiber Polarity



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




Tan pull-tab denotes multimode optics

Blue label denotes 50-meter reach

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)



Note

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

Transceiver Back-Shell Label Serial Number Legend

Symb ol	Meaning	Notes
MT	Manufacturer name (Mellanox Technologies)	2 digits (alphanumeric)
YY	Year of manufacturing	2 last digits of the year (numeric)
WW	Week of manufacturing	2 digits (numeric)
JC <u>or</u> DM	Manufacturer Site: JC – Option 1 (China) DM – Option 2 (Malaysia)	Two characters
SSSS S	Serial number	5 digits (decimal numeric) for serial number, starting from 00001.

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

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.

Connector and Cabling Details

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



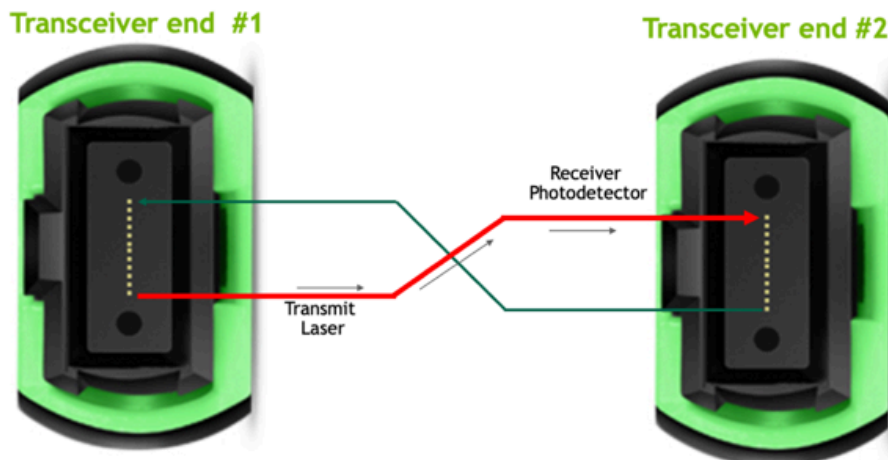
NVIDIA Supplied Crossover Type-B Fiber Cables

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



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

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.



Cable Management Guidelines

For more information and general interconnect management and installation, see [NVIDIA Cable Management Guidelines and FAQ Application Note](#).


Part Numbers and Description

OPN	Description
MMA4Z00-NS	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, finned
MMA4Z00-NS-FLT	NVIDIA twin port transceiver, 800(2x400)Gbps, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, flat top

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

OPN	4-channel MPO/APC to Two 2-channel MPO/APC
MFP7E20-N005	5m
MFP7E20-N007	7m
MFP7E20-N010	10m
MFP7E20-N015	15m
MFP7E20-N020	20m
MFP7E20-N030	30m
MFP7E20-N050	50m

Document Revision History

Rev	Date	Description
1.6	Feb. 2025	Updated Flat-top module mechanical drawing
1.5	Apr. 2024	Added Flat-top head detail mechanical drawing
1.4	Jan. 2024	Updated low power mode output.
1.3	Dec. 2023	Updated Regulatory Compliance section
1.2	Nov. 2023	Added mechanical drawings. Updated specifications.
1.1	Mar. 2023	- Updated the document for Ethernet support. - Minor text edits.
1.0	Dec. 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. NVIDIA Corporation (“NVIDIA”) makes no 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2024, NVIDIA. PDF Generated on 05/05/2025