# MMA1B00-E100 100Gb/s Transceiver

# Product Specifications

# Table of Contents

# 1 Introduction

NVIDIA® MMA1B00-E100 pluggable optical transceiver is designed for use in 100Gb/s InfiniBand link protocol applications.

This SFF-8665 compliant transceiver is a flexible alternative to an Active Optical Cable (AOC), as it combines high port density and configurability with longer reach than passive copper cables in the data centers. The MMA1B00 transceiver has a standard QSFP28 port on the electrical side towards the host system.

The optical cable is attached via a standard MPO connector. The optical interface comprises four optical beams/fibers in each direction, intended for multimode optical cable attachments. Each fiber/beam operates at signaling rates up to 25.78125GBd.

> ⚠ NVIDIA EDR AOCs and transceivers including MMA1B00-E100 work in NVIDIA end-to-end Ethernet applications. However, the Ethernet version of the transceiver (MMA1B00-C100D) is generally not supported in InfiniBand applications.

Rigorous production testing ensures the best out-of-the-box installation experience, performance, and durability.

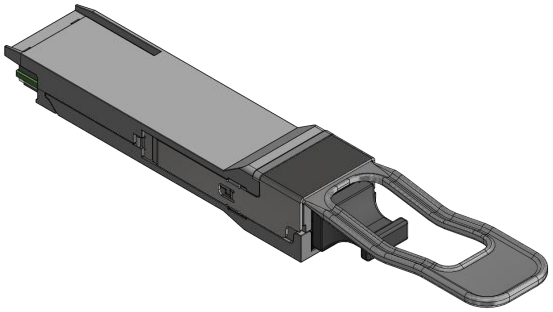

> ⚠ Images are for illustration purposes only. Product labels and colors may vary.

## 1.1 Key Features

- Up to 100Gb/s
- Programmable Rx Output Amplitude
- Programmable Rx Output Emphasis
- Programmable Tx Input Equalizer
- Selectable retiming
- SFF-8665 compliant
- 3.3V power supply
- QSFP28 Power Class 3
- Class 1 Laser Safety
- Maximum link length of 70m on OM3, and 100m on OM4 multimode fiber at 100Gb/s
- Hot pluggable
- RoHS compliant
- IEEE 802.3 100GBASE-SR4 compliant
- $I^2C$ management interface

# 2 Pin Description

## 2.1 SFP28 Pin Function Description

The transceiver's pin assignment is SFF-8679 compliant.

| 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 | Grounds |
| 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 |

Notes:

- GND is the symbol for signal and supply (power) common for the QSFP28 module. All are common within the module and all module voltages are referenced to this potential unless otherwise noted. Connect these directly to the host board signal common ground plane.
- VccRx, Vcc1 and VccTx are the receiving and transmission power suppliers and shall be applied concurrently.

- Vcc Rx, Vcc1 and Vcc Tx may be internally connected within the module in any combination. The connector pins are each rated for a maximum current of 1000 mA.

### 2.1.1 QSFP28 Module Pad Layout

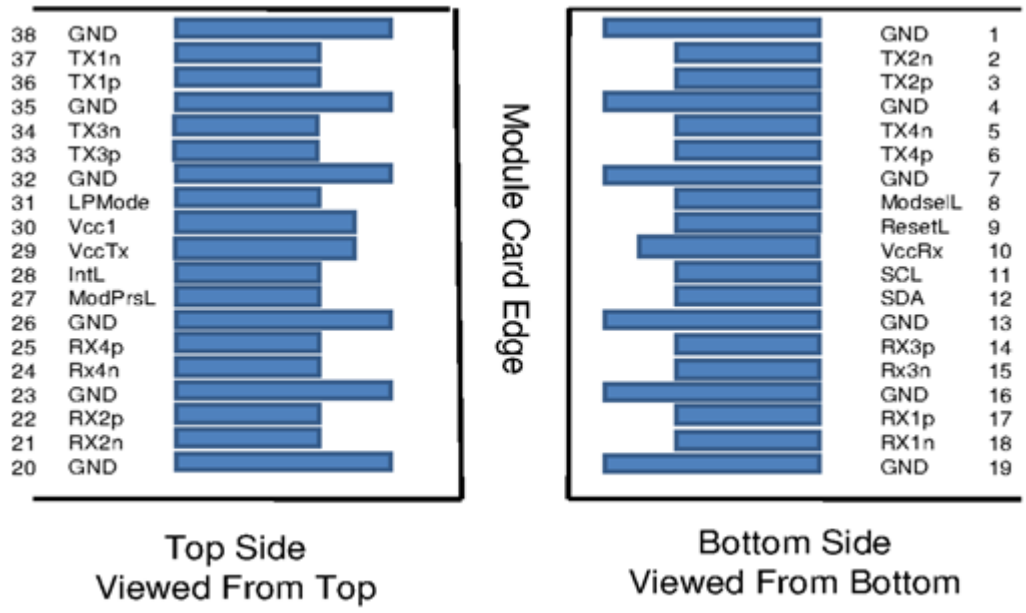


## 2.2 Control Signals

This transceiver is SFF-8636 compliant. This means that the control signals shown in the pad layout support the following functions:

| Name | Function | Description |
|---|---|---|
| ModPrsL | Output, asserted low | 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 | Input, asserted Low | 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 | Input, asserted Low | 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. |

| Name | Function | Description |
|---|---|---|
| LPMode | Input, asserted high | 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 | Output, asserted low | 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’. |

The low-speed signals are Low Voltage TTL (LVTTL) compliant (except for SCL and SDA signals).

## 2.3 Diagnostics and Other Features

The transceiver complies with the SFF-8665 specification and has the following key features:

Physical layer link optimization:

- Programmable Tx input equalization
- Programmable Rx output amplitude
- Programmable Rx output pre-emphasis
- Tx/Rx CDR control
  by default enabled for 100 GbE operation, disable it for 40G operation

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.

# 3 Specifications

## 3.1 Absolute Maximum Specifications

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

| Parameter | Min | Max | Units |
|---|---|---|---|
| Supply voltage | -0.3 | 3.465 | V |
| Data input voltage | -0.3 | 3.465 | V |
| Control input voltage | -0.3 | 4 | V |
| Damage Threshold, per optical lane | 3.4 | --- | dBm |
| Storage temperature | -40 | 85 | °C |

## 3.2 Recommended Operating Conditions and Power Supply Requirements

| Parameter | Min | Typ | Max | Units |
|---|---|---|---|---|
| Supply voltage | 3.135 | 3.3 | 3.465 | V |
| Power consumption (no retiming) | --- | 1.5 | 1.8 | W |
| Power consumption (retiming on all lanes) | --- | 2.2 | 2.5 | W |
| Supply noise tolerance (10 Hz – 10 MHz) | 66 | --- | --- | mVpp |
| Operating case temperature | 0 | --- | 70 | °C |
| Operating relative humidity | 5 | --- | 85 | % |

## 3.3 Electrical Specifications

| Parameter | Min | Typ | Max | Units |
|---|---|---|---|---|
| Differential input return loss | Per IB requirements | | | dB |
| Differential output return loss | | | | dB |
| Common mode output return loss | | | | dB |
| J2 | --- | --- | J2 | --- |
| J9 | --- | --- | J9 | --- |
| Output transition time | 10 | --- | --- | ps |

## 3.4 Optical Specifications

| Parameter | Min | Typ | Max | Units |
|---|---|---|---|---|
| **Transmitter (per lane)** | | | | |
| Signaling Speed (with retiming) | 25.78025 | 25.78125 | 25.78225 | GBd |
| Signaling Speed (without retiming) | 0.300 | --- | 25.78125 | GBd |
| Center Wavelength | 840 | --- | 860 | nm |
| Average Launch Power | -8.4 | --- | 2.4 | dBm |
| Transmit OMA | -6.4 | --- | 3 | dBm |
| Extinction Ratio | 2 | --- | --- | dB |
| Transmitter and Dispersion Eye Closure | --- | --- | 4.3 | dB |
| Average Launch power at Tx Squelch | | | -30 | dBm |
| Eye Crossing | 45 | --- | 55 | % |
| Optical return loss tolerance | --- | --- | 12 | dB |
| Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}<br>Hit ratio 1.5 x $10^{-3}$ hits per sample | {0.3, 0.38, 0.45, 0.35, 0.41, 0.5} | | | --- |
| Reach on OM3 multi-mode fiber | 70 | --- | --- | m |
| Reach on OM4 multi-mode fiber | 100 | --- | --- | m |
| **Receiver (per lane)** | | | | |
| Signaling Speed (with retiming) | 25.78025 | 25.78125 | 25.78225 | GBd |
| Signaling Speed (without retiming) | 0.300 | --- | 25.78125 | GBd |
| Center Wavelength | 840 | --- | 860 | nm |
| Damage Threshold (see note 1) | 3.4 | --- | --- | dBm |
| Receiver Reflectance | --- | --- | -12 | dB |
| LOS De-Assert | --- | --- | -12 | dBm |
| LOS Assert | -30 | --- | --- | dBm |
| LOS Hysteresis | 0.5 | --- | --- | dB |
| Unstressed Receiver Sensitivity (OMA) at BER 1E-12 (see note 2) | --- | --- | -6 | dBm |
| Stressed Receiver Sensitivity (OMA), each lane (see note 3) | --- | --- | -5.2 | dBm |
| Conditions for Stressed Receiver Sensitivity (see note 4) | See below | | | |
| **Test Conditions** | | | | |
| Stressed eye closure (SEC), lane under test | 4.3 | | | dB |
| Stressed eye J2 Jitter, lane under test | 0.39 | | | UI |

| Parameter | Min | Typ | Max | Units |
|---|---|---|---|---|
| Stressed eye J4 Jitter, lane under test (max) | 0.53 | | | UI |
| OMA of each aggressor lane | 3 | | | dBm |
| Stressed receiver eye mask definition {X1, X2, X3, Y1, Y2, Y3}<br>Hit ratio 5 x $10^{-5}$ hits per sample | {0.28, 0.5, 0.5, 0.33, 0.33, 0.4} | | | --- |

Notes:

1. The receiver may not operate correctly at this input power level.
2. All Tx channels on and all Rx channels on with AOP 3dB greater than the tested channel. Injected optical eye must comply with the transmitter eye mask definition.
3. Measured with conformance test signal at TP3 (ref. IEEE 802.3 100GBASE-SR4) for BER < 5E-5.
4. The test conditions are for measuring stressed receiver sensitivity only – not characteristics of the receiver.

## 3.5 Electrostatic Discharge (ESD)

This product is compatible with ESD levels in typical data center operating environments and certified in accordance with the standards listed in the Regulatory Compliance Section. The product is shipped with protective caps on all connectors to protect it during shipping. In normal handling and operation of high-speed cables and optical transceivers, ESD is of concern during insertion into the OSFP cage of the server/switch. Hence, standard ESD handling precautions must be observed. These include use of grounded wrist/shoe straps and ESD floor wherever a cable/transceiver is extracted/inserted. Electrostatic discharges to the exterior of the host equipment chassis after installation are subject to system level ESD requirements.

## 3.6 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.

Prior to insertion of the fiber cable, clean the cable connector to prevent contamination from it. The dust cap ensures that the optics remain clean and no additional cleaning should be needed. In the event of contamination, standard cleaning tools and methods should be used. Liquids must not be applied.

## 3.7 Mechanical Specifications

### 3.7.1 Mechanical Dimensions

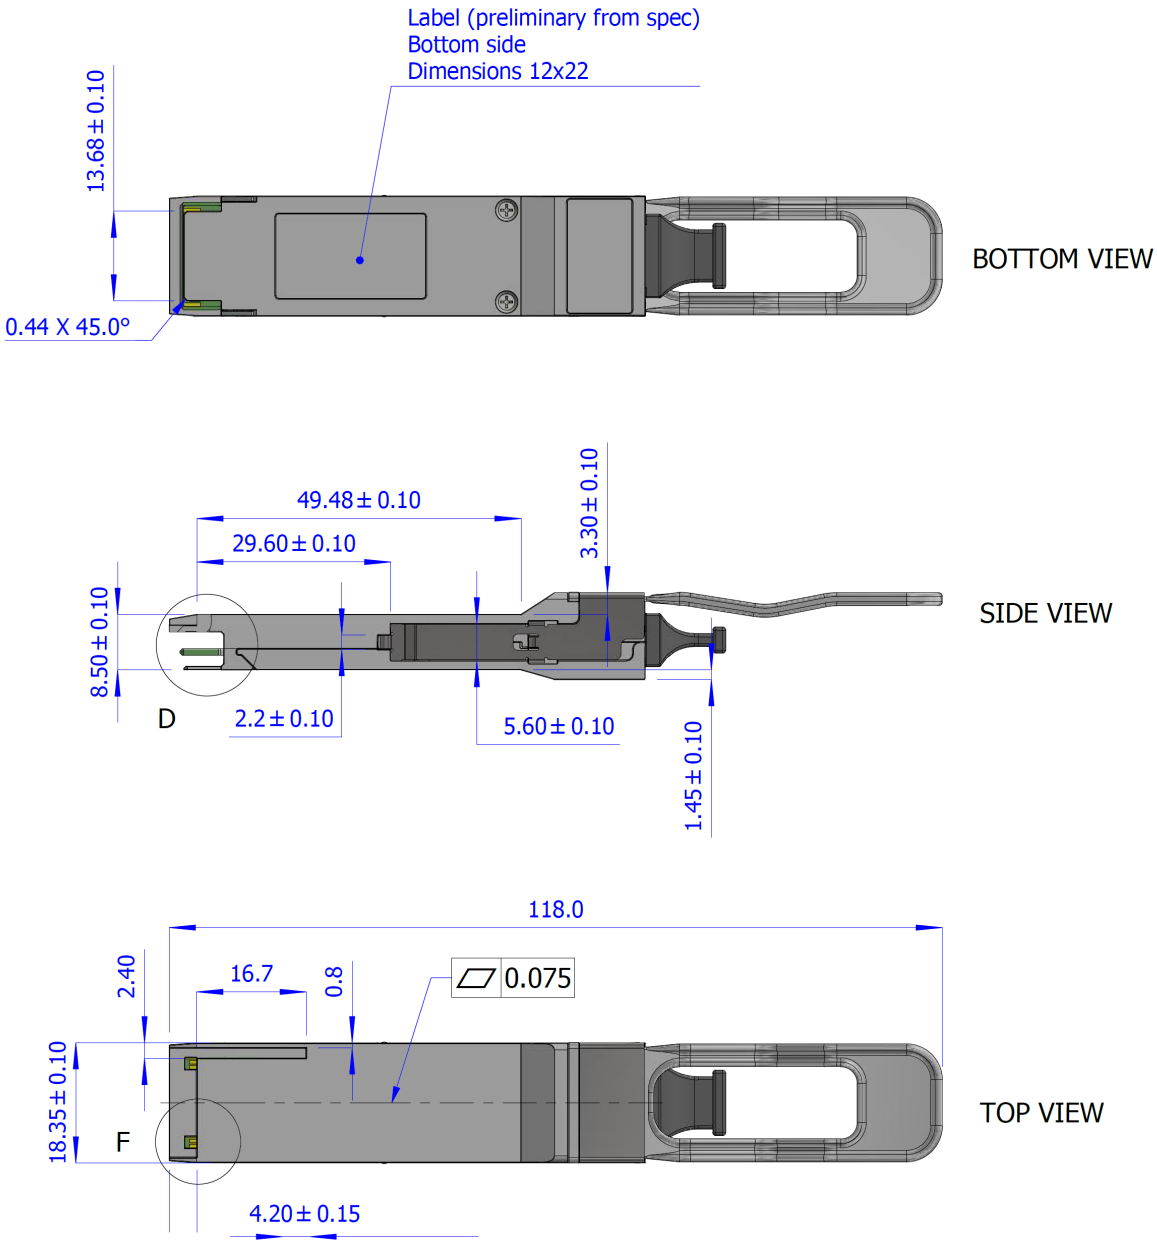

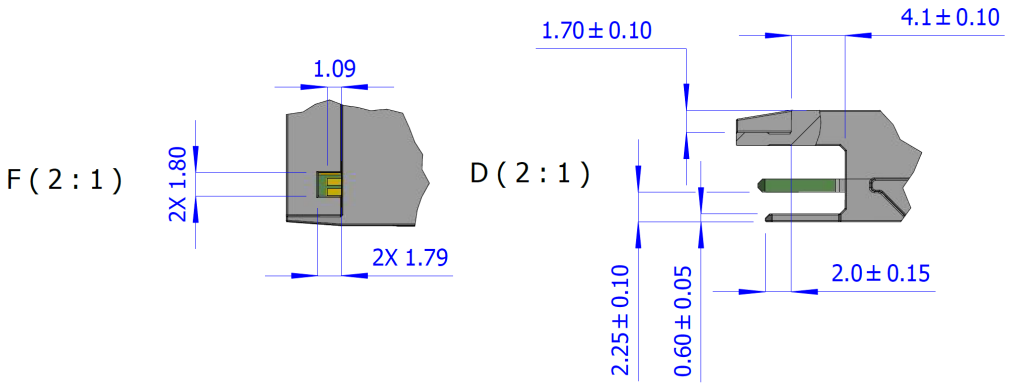


## 3.8 Memory Map

The transceiver's memory map is compliant with the QSFP Management interface specification SFF-8636. See also the NVIDIA LinkX® Memory Map Application Note (MLNX-15-5926).

## 3.9 Label

The following label is applied on the transceiver's backshell:

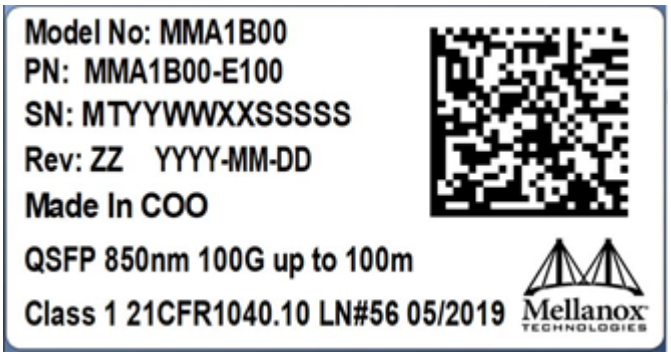


(sample illustration)

### 3.9.1 Backshell Label Legend

| Symbol | Meaning | Notes |
|---|---|---|
| SN - Serial Number | | |
| MT | Manufacturer name | 2 characters (MT) |
| YY | Year of manufacturing | 2 digits |
| WW | Week of manufacturing | 2 digits |
| DM | Manufacturer site | 2 characters |
| ZZZZZ | Serial number | 5 digits for serial number, starting from 00001. Reset at start of week to 00001. |
| **Miscellaneous** | | |
| ZZ | HW and SW revision | 2 alpha-numeric characters |
| YYYY | Year of manufacturing | 4 digits |

| Symbol | Meaning | Notes |
|---|---|---|
| MM | Month of manufacturing | 2 digits |
| DD | Day of manufacturing | 2 digits |
| COO | Country of origin | E.g. China or Malaysia |
|  | Quick response code | Serial number (MTYYWWXXSSSSS) |

## 3.10 Regulatory Compliance and Classification

The laser module is classified as class I according to IEC 60825-1, IEC 60825-2 and 21 CFR 1040 (CDRH).

- Safety: FDA/CDRH, TUV, UL/CSA, ACMA
- EMC: NTS

Ask your NVIDIA FAE for a zip file of the certifications for this product.

## 3.11 FCC Class A Notice

Each of the devices complies with CFR47 FCC Class A Part 15 of the FCC Rules. Operation is subject to the following two conditions:

1. This device may not cause harmful interference.
2. This device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur during installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/television technician for help.

Modifications: Any modifications made to this device that are not approved by NVIDIA may void the authority granted to the user by the FCC to operate this equipment.

# 4 Ordering Information

| Ordering Part Number | Description |
|---|---|
| MMA1B00-E100 | Transceiver, IB EDR, up to 100Gb/s, QSFP28, MPO, 850nm, up to 100m |

# 5 References

1. NVIDIA LinkX® SFF Memory Map Application Note
2. NVIDIA LinkX® Measuring Eye Parameters Application Note
3. NVIDIA LinkX® Racking and Cabling Guidelines
4. NVIDIA LinkX® FAQ

For further information, please contact your sales representative or the NVIDIA support team.

# 6 Appendix: Optical Connector and Patch Cable

## 6.1 Optical Connector

The optical port in the parallel 2x 4-lane optical transceiver is a male MPO connector with alignment pins, mating with fiber-optic cables with female MPO connector. The connector contains a 12-channel MT ferrule.

*QSFP28 Optical Male MPO Connector of the Transceiver (front view)*

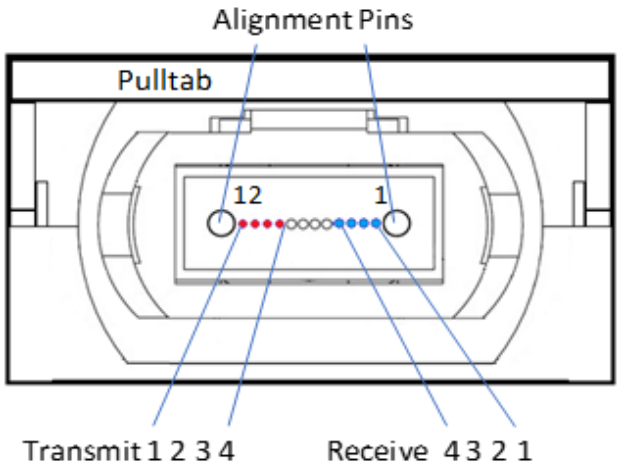


*Female MPO Cable Connector*

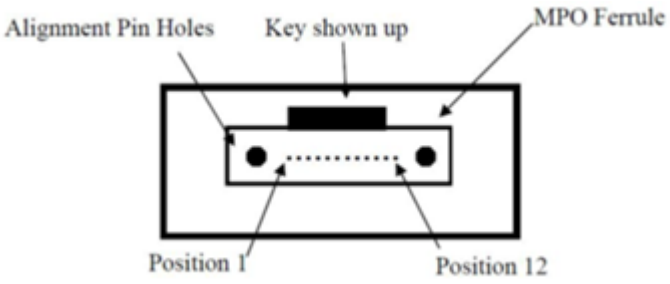


Reference: IEC specification IEC 61754-7.\

## 6.2 Patch Cable

The fiber which connects with Transceiver A's lane 1 must end at Transceiver B's lane 12 at the far end of the cable. For example, position 1 of MPO connector A at the near end of the cable connects to position 12 of the far end MPO connector B. This calls for a 'crossed' MPO cable, which is often referred to as patch cable type B. See interconnect scheme below.

*MPO to MPO Patch Cable Fiber Connections*

| Connector A | Connection | Connector B |
|---|---|---|
| 1 | <-- | 12 |

| Connector A | Connection | Connector B |
|---|---|---|
| 2 | <-- | 11 |
| 3 | <-- | 10 |
| 4 | <-- | 9 |
| 5 | Not Connected | 8 |
| 6 | Not Connected | 7 |
| 7 | Not Connected | 6 |
| 8 | Not Connected | 5 |
| 9 | --> | 4 |
| 10 | --> | 3 |
| 11 | --> | 2 |
| 12 | --> | 1 |

As shown in the above table, the 4 middle positions are not in use, indicating that the cable may have 8 or 12 fibers. Only 8 fibers are used.

The fiber is a standard OM3 or OM4 multi-mode fiber. The maximum length is detailed in Optical Specifications.

*MPO Fiber Patch Cables OM3 (left) and OM4 (right)*

OM4 cables may have aqua or pink colored connectors, depending on the make and model.

NVIDIA does not provide MPO patch cables. Multiple cable vendors provide these patch cables.

Multiple MPO patch cables can be connected in a series, e.g. via Optical Distribution Frames (ODFs). Each added connector pair increases modal dispersion and reflections in the link which impairs performance. An odd number of ‘crosses’ must be used between transceivers at the two ends.

# 7 Document Revision History

| Version | Date | Description of Change |
|---|---|---|
| 1.6 | Apr. 2024 | Updated product mechanical drawings and graphics. |
| 1.5 | Nov. 2021 | Reformatted and rebranded; migrated to HTML. |
| 1.4 | Apr. 2021 | Removed “BER better than $10^{-15}$” from Key Features.<br>Added Electrostatic Discharge (ESD) and Handling and Cleaning sections.<br>Updated Transmit OMA, Extinction Ratio, and Transmitter and dispersion eye closure values in the Optical Specifications table.<br>Updated document template, legal notice, and introduction. |
| 1.3 | Nov. 2020 | Added: Control Signals description, optical connector interface information in Appendix A and added References section.<br>Updated DDM info under Diagnostics and Other Features and the Regulatory Compliance and Classification list.Minor text edits. |
| 1.2 | Jul. 2020 | Converted to new template.<br>Added information regarding Ethernet version of the module to the introduction. |
| 1.1 | Aug. 2016 | Updated Optical Specifications<br>Updated minimum temperature storage |
| 1.0 | Mar. 2016 | 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/