

Arista 400G Transceivers and Cables: Q&A

What are the benefits of moving to 400G technology?

Arista's 400G platforms allow data centers and high-performance computing environments to address growing needs for higher bandwidth at lower cost and power per gigabit. Key benefits include:

- Increase switching bandwidth by a factor of 4. Migrating from 100G to 400G systems increases the bandwidth per RU from 3.2-3.6T to 12.8-14.4T / RU.
- Enable higher density 100G ports using optical or copper breakouts. A 32 port 1RU 400G system enables 128 100GE ports / RU. This enables a single Top of Rack (TOR) leaf switch to connect to multiple racks of servers or Network Interface Cards (NICs).
- Reduce the number of optical fiber links, connectors, and patch panels by a factor of 4 when compared to 100G platforms for the same aggregate bandwidth.
- Enable 2-4X lower cost and power / bit, reducing capex and opex.

What form-factors are used for 400G optics and cables?

There are two form-factors defined for 400G optical modules and cables:

- **The OSFP:** The OSFP stands for "Octal Small Form-factor Pluggable". It is described as an "Octal" module because the electrical interface of an OSFP connector consists of 8 electrical lanes, running at 50Gb/s each, for a total of bandwidth of 400Gb/s.
- **The QSFP-DD:** The QSFP-DD stands for "Quad Small Form-factor Pluggable (QSFP) – Double Density (DD)". The electrical interface of a QSFP-DD connector also has 8 electrical lanes, running at 50Gb/s each, for a total bandwidth of 400Gb/s. The QSFP-DD form factor is similar to the QSFP form factor, except a second row of electrical contacts has been added to increase the number of high-speed electrical lanes from 4 (in a QSFP) to 8 (in a QSFP-DD).

What 400 form-factors does Arista support?

Arista supports both OSFP and QSFP-DD form-factors – i.e. Arista 400G platforms will be offered in both OSFP and QSFP-DD variants (a detailed comparison of OSFP and QSFP-DD form factors is contained later in this document).

Can I plug an OSFP module into a QSFP-DD port, or a QSFP-DD module into an OSFP port?

No. The OSFP and the QSFP-DD are two physically distinct form factors. If you have an OSFP system, then OSFP optics and cables must be used. If you have a QSFP-DD system, then QSFP-DD optics and cables must be used.

Can I have an OSFP on one end of a 400G link, and a QSFP-DD on the other end?

Yes. The OSFP and QSFP-DD describe the physical form factor of the module. As long as the Ethernet media types are the same (i.e. both ends of the link are 400G-DR4, or 400G-FR4 etc.), OSFP and QSFP-DD modules will interop with each other.

What 400G Transceivers and Cables are available from Arista?

Arista supports a full range of 400G optical transceivers, Active Optical Cables (AOCs) and Direct Attach Copper cables (DACs) in both OSFP and QSFP-DD form factors. The tables below summarize the 400G connectivity options available from Arista.

400G Optical Transceivers		
OSFP Part No.	QSFP-DD Part No.	Product Description
OSFP-400G-SR8	QDD-400G-SR8	400GBASE-SR8 transceiver, MPO-16 APC connector, up to 100m over parallel multi-mode OM4 fiber (MMF). The 400G-SR8 can breakout to 2x 200G-SR4 or 2x 100G-SR4 optics.
OSFP-400G-SR8-C	QDD-400G-SR8-C	400GBASE-SR8 and 8x 50G-SR / 25G-SR OSFP transceiver, MPO-16 APC connector, up to 100m over parallel OM4 multi-mode fiber. 400G-SR8-C transceivers include all the functionality of the 400G-SR8 transceiver, with the additional capability to breakout and connect to 8x 50G-SR or 8x25G-SR optics ("Channelized" configuration).
OSFP-400G-SRBD	QDD-400G-SRBD	400GBASE-SR4.2 (400G-BIDI), MPO-12 connector, up to 70m / 100m over parallel OM3 / OM4 MMF. The 400G-SRBD is software configurable for point-to-point 400GE links (400GBASE-SR4.2), breakout to 4x 100GBASE-BIDI (100G-SRBD) or breakout to 4x 100GBASE-SR1.2 optics.
OSFP-400G-VSR4	QDD-400G-VSR4	400GBASE-VSR4 Transceiver, MPO-12 APC connector, up to 50m over parallel multi-mode OM4 fiber (MMF).
OSFP-400G-DR4	QDD-400G-DR4	400GBASE-DR4 transceiver, up to 500m over parallel single mode fiber (SMF). The 400G-DR4 can break out to 4 x 100G, and interop with 4 x 100G-DR QSFPs.
OSFP-400G-XDR4	QDD-400G-XDR4	400GBASE-XDR4 "eXtended reach DR4" transceiver, up to 2km over parallel SMF. The 400G-XDR4 can break out to 4 x 100G, and interop with 4 x 100G-FR QSFPs.
OSFP-400G-PLR4	QDD-400G-PLR4	400GBASE-PLR4 "Parallel Long Reach" transceiver, up to 10km over parallel SMF. The 400G- PLR 4 can break out to 4 x 100G, and interop with 4 x 100G-LR QSFPs.
OSFP-400G-FR4	QDD-400G-FR4	400GBASE-FR4 transceiver, up to 2km over duplex SMF.
OSFP-400G-LR4	QDD-400G-LR4	400GBASE-LR4 transceiver, up to 10km over duplex SMF.
OSFP-400G-2FR4	-	2 x 200GBASE-FR4 transceiver, up to 2km over two pairs of duplex SMF. The 2FR4 transceiver can break out to 2 x 200G, and interop with 2 x 200G-FR4 QSFP transceivers.

400G AOC (Active Optical Cables)		
OSFP Part Number	QSFP-DD Part No.	Product Description
AOC-O-O-400G-1M	AOC-D-D-400G-1M	400G AOC, 1 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-3M	AOC-D-D-400G-3M	400G AOC, 3 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-5M	AOC-D-D-400G-5M	400G AOC, 5 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-7M	AOC-D-D-400G-7M	400G AOC, 7 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-10M	AOC-D-D-400G-10M	400G AOC, 10 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-15M	AOC-D-D-400G-15M	400G AOC, 15 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-20M	AOC-D-D-400G-20M	400G AOC, 20 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-25M	AOC-D-D-400G-25M	400G AOC, 25 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
AOC-O-O-400G-30M	AOC-D-D-400G-30M	400G AOC, 30 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)

400G 'Straight Through' Passive Copper Cables (or Direct Attach Cables or 'DACs')		
OSFP Part Number	QSFP-DD Part No.	Product Description
CAB-O-O-400G-1M	CAB-D-D-400G-1M	400G Passive DAC, 1 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
CAB-O-O-400G-2M	CAB-D-D-400G-2M	400G Passive DAC, 2 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)
CAB-O-O-400G-3M	CAB-D-D-400-2.5 & CAB-D-D-400-3M	400G Passive DAC, 3 / 2.5 meter (OSFP to OSFP, or QSFP-DD to QSFP-DD)

400G Breakout Passive DACs		
OSFP Part Number	QSFP-DD Part No.	Product Description
OSFP or QSFP-DD to 2x 200G QSFPs (with each 200G QSFP using 4 x 50G PAM-4 electrical lanes)		
CAB-O-2Q-400G-1M	CAB-D-2Q-400G-1M	400G Passive DAC, 1 meter (OSFP or QSFP-DD to 2x 200G-4 QSFP)
CAB-O-2Q-400G-2M	CAB-D-2Q-400G-2M	400G Passive DAC, 2 meter (OSFP or QSFP-DD to 2x 200G-4 QSFP)
CAB-O-2Q-400G-3M	CAB-D-2Q-400-2.5 & CAB-D-2Q-400-3M	400G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 2x 200G-4 QSFP)
OSFP or QSFP-DD to 4x 100G QSFPs (with each 100G QSFP using 2 x 50G PAM-4 electrical lanes)		
CAB-O-4Q-400G-1M	CAB-D-4Q-400G-1M	400G Passive DAC, 1 meter (OSFP or QSFP-DD to 4x 100G-2 QSFP with 2 x 50G lanes per QSFP)
CAB-O-4Q-400G-2M	CAB-D-4Q-400G-2M	400G Passive DAC, 2 meter (OSFP or QSFP-DD to 4x 100G-2 QSFP with 2 x 50G lanes per QSFP)
CAB-O-4Q-400G-3M	CAB-D-4Q-400-2.5 & CAB-D-4Q-400-3M	400G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 4x 100G-2 QSFP with 2 x 50G lanes per QSFP)
OSFP or QSFP-DD to 8x 50G SFP (with each 50G QSFP using a 1 x 50G PAM-4 electrical lane)		

CAB-O-8S-400G-1M	CAB-D-8S-400G-1M	400G Passive DAC, 1 meter (OSFP or QSFP-DD to 8x 50G SFP)
CAB-O-8S-400G-2M	CAB-D-8S-400G-2M	400G Passive DAC, 2 meter (OSFP or QSFP-DD to 8x 50G SFP)
CAB-O-8S-400G-3M	CAB-D-8S-400-2.5	400G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 8x 50G SFP)

400G Breakout Active Copper Cables (ACCs)		
OSFP Part Number	QSFP-DD Part No.	Product Description
OSFP or QSFP-DD to 4x 100G QSFPs (with each 100G QSFP using 4 x 25G NRZ electrical lanes)		
H-O400-4Q100-1M	H-D400-4Q100-1M	400G Active Copper Cable, 1 meter (OSFP or QSFP-DD to 4x 100G-4 QSFP)
H-O400-4Q100-2M	H-D400-4Q100-2M	400G Active Copper Cable, 2 meters (OSFP or QSFP-DD to 4x 100G-4 QSFP)
H-O400-4Q100-3M	H-D400-4Q100-3M	400G Active Copper Cable, 3 meters (OSFP or QSFP-DD to 4x 100G-4 QSFP)
H-O400-4Q100-5M	H-D400-4Q100-5M	400G Active Copper Cable, 5 meters (OSFP or QSFP-DD to 4x 100G-4 QSFP)

200G Breakout DACs for connection to legacy 100G / 50G QSFP ports and 25G SFP ports (Note: The OSFP / QSFP-DD port must be operated at 200G to allow breakouts with 25G NRZ electrical lanes)		
OSFP Part Number	QSFP-DD Part No.	Product Description
OSFP or QSFP-DD to 2x 100G QSFPs (with each 100G QSFP using 4 x 25G NRZ electrical lanes)		
CAB-O-2Q-200G-1M	CAB-D-2Q-200G-1M	200G Passive DAC, 1 meter (OSFP or QSFP-DD to 2x 100G-4 QSFP)
CAB-O-2Q-200G-2M	CAB-D-2Q-200G-2M	200G Passive DAC, 2 meter (OSFP or QSFP-DD to 2x 100G-4 QSFP)
CAB-O-2Q-200G-3M	CAB-D-2Q-200-2.5	200G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 2x 100G-4 QSFP)
OSFP or QSFP-DD to 4x 50G QSFPs (with each 50G QSFP using 2 x 25G NRZ electrical lanes)		
CAB-O-4Q-200G-1M	CAB-D-2Q-200G-1M	200G Passive DAC, 1 meter (OSFP or QSFP-DD to 4x 50G-2 QSFP with 2 x 25G lanes per QSFP)
CAB-O-4Q-200G-2M	CAB-D-2Q-200G-2M	200G Passive DAC, 2 meter (OSFP or QSFP-DD to 4x 50G-2 QSFP with 2 x 25G lanes per QSFP)
CAB-O-4Q-200G-3M	CAB-D-2Q-400-2.5	200G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 4x 50G-2 QSFP with 2 x 25G lanes per QSFP)
OSFP or QSFP-DD to 8x 25G SFP (each 25G QSFP uses a 1 x 25G NRZ electrical lane)		
CAB-O-8S-200G-1M	CAB-D-8S-200G-1M	200G Passive DAC, 1 meter (OSFP or QSFP-DD to 8x 25G SFP)
CAB-O-8S-200G-2M	CAB-D-8S-200G-2M	200G Passive DAC, 2 meter (OSFP or QSFP-DD to 8x 25G SFP)
CAB-O-8S-200G-3M	CAB-D-8S-200-2.5	200G Passive DAC, 3 / 2.5 meter (OSFP or QSFP-DD to 8x 25G SFP)

What is the difference between a '400G' and '200G' breakout DAC?

Arista offers passive copper breakout cables in 400G and 200G versions. Both versions include breakouts from an OSFP / QSFP-DD port to 2x QSFP, 4x QSFP, and 8x SFP ports. The key differences between the 400G and 200G versions are:

1. The 400G cables are designed to carry 8x 50Gb/s PAM-4 electrical lanes, while the 200G cables are designed to carry 8x 25Gb/s NRZ electrical lanes (see a later section of this document for a detailed explanation of PAM-4 and NRZ signaling). The high-speed 50G PAM-4 signals place more stringent signal integrity requirements on the DACs, and the 400G cables are specified to support these signals. The 400G cables will also support lower speed 25Gb/s NRZ signals, but there may be interop issues when using 400G cables in legacy systems designed for 25Gb/s NRZ cables (see explanation below).
2. Each connector on the DAC cable (i.e. the OSFP / QSFP-DD and QSFP / SFP connectors) contains an EEPROM which is programmed to indicate the signal format and max data rate supported by the cable. The EEPROMs in the 400G cables indicate support for up to 50G PAM-4 data rates, while the EEPROMs in the 200G cables indicate support up to only 25G NRZ data rates.

Breakout DACs are often used to connect Arista switches to other non-Arista equipment such as server NICs (Network Interface Cards). The NIC software will usually read the DAC connector's EEPROM to determine what type of device has been plugged in. The software installed on many legacy NICs may not recognize the newer signal speeds and modulation formats, and hence may not successfully recognize cables programmed for 50Gb/s PAM-4 signals. The 200G cables can be used when connecting to legacy equipment that will only recognize DACs programmed for 25Gb/s data rates.


Can I plug a QSFP module into an OSFP port?

Yes. A QSFP (40G or 100G) module can be inserted into an OSFP port by using a simple, completely passive, mechanical adapter. This adapter is available from Arista, with part number ADPT-O-Q-100G. The adapter, and the adapter with a QSFP plugged into it, are shown below.

OSFP to QSFP adapter (ADPT-O-Q-100G)



OSFP to QSFP adapter, with a 100G QSFP plugged in



When using a QSFP module in an OSFP port, the OSFP port must be configured for a data rate of 100G (or 40G), instead of 400G.

Can I plug a 100G QSFP module into a QSFP-DD port?

Yes. A QSFP (40G or 100G) module can also be inserted into a QSFP-DD port (without a mechanical adapter). When using a QSFP module in a QSFP-DD port, the QSFP-DD port must be configured for a data rate of 100G (or 40G), instead of 400G.

What do the suffixes “SR8, SR8-C, SRBD, VSR4, DR4 / XDR4 / PLR4, FR4 / LR4 and 2FR4” stand for?



In general, the letters refer to reach or optics technology, and the number refers to the number of optical channels:

- i) **SR8:** “SR” refers to 100m reach using multi-mode fiber, and “8” implies there are 8 optical channels. Each of the 8 optical channels from an SR8 module are carried on separate fibers, resulting in a total of 16 fibers (8 Tx and 8 Rx). Each optical channel operates at 50Gb/s. The SR8 module uses an MPO-16 APC connector to connect to 8 fiber pairs. The SR8 module can be configured for 400G-SR8, 2x 200G-SR4 and (operating at ½ rate) 2x 100G-SR4.
- ii) **SR8-C:** Arista’s 400G-SR8-C transceivers include all the functionality of the 400G-SR8 transceiver, plus the ability to be configured as 8x 50G-SR or 8x25G-SR (“Channelized” configuration)
- iii) **SRBD:** “SR” refers to 100m reach using multimode fiber, and “BD” refers to Bi-directional (BIDI) optical technology. 400G-SRBD transceivers are also referred to as “400G-BIDI” or “400G-SR4.2” transceivers. Arista’s 400G-SRBD transceivers are compliant with IEEE 400GBASE-SR4.2 and are software configurable to breakout into either 4x 100G-BIDI (QSFP-100G-SRBD) or 4x 100G-SR1.2 links. Refer to later sections in this document for more information on the 400G-SRBD.
- iv) **VSR4:** “VSR” refers to 50m reach using multimode fiber, and the “4” implies there are 4 optical channels. Each of the 4 optical channels are carried on separate fibers, resulting in a total of 4 pairs of fibers. Each optical channel operates at 100Gb/s.

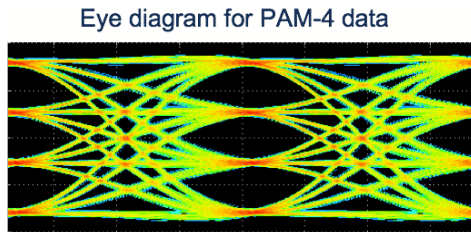
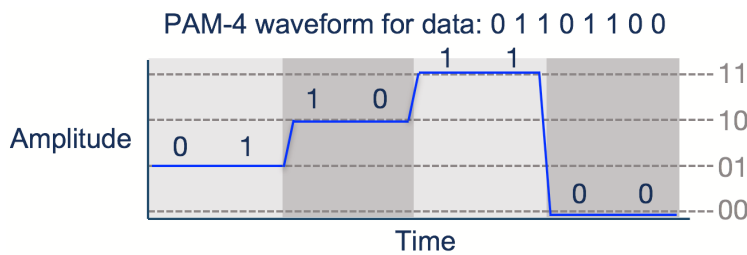
There are two IEEE defined standards for 100G/wave MMF optics: 400GBASE-SR4 (for 100m reach over parallel OM4 MMF), and 400GBASE-VR4 (for 50m reach over parallel OM4 MMF). Arista’s 400G-VSR4 transceivers are fully compliant with the 400GBASE-VR4 standard, and optically interoperable with both 400GBASE-SR4 and 400GBASE-VR4 transceivers over a reach of 50m.
- v) **DR4, XDR4 and PLR4:** “DR” / “XDR” / “LR” refer to 500m / 2km / 10km reach using single-mode fiber, and “4” implies there are 4 optical channels. Each of the 4 optical channels are carried on separate fibers, resulting in a total of 4 pairs of fibers. Each optical channel operates at 100Gb/s. The DR4 / XDR4 / PLR4 modules use the widely deployed MPO-12 APC SMF connector to connect to 4 fiber pairs.
- vi) **FR4 / LR4:** “FR” / “LR” refers to 2km / 10km reach using single-mode fiber, and “4” implies there are 4 optical channels. Unlike the DR4 and SR8, all 4 optical channels from an FR4 / LR4 are multiplexed onto one fiber, resulting in a total of 2 fibers from the module (1 Tx and 1 Rx). Each optical channel operates at 100Gb/s. The modules use a duplex LC optical connector.
- vii) **2FR4:** The “2FR4” refers to 2 x 200G-FR4 links, and has a reach of 2km using single-mode fiber. Each of the 200G FR4 links has 4 optical channels, multiplexed onto one fiber pair (1 Tx and 1 Rx per 200G link). A 400G-2FR4 module has 2 of these links, resulting in a total of two pairs of single mode fiber (or 4 fibers total), and a total of 8 optical channels. Each optical channel operates at 50Gb/s. The 2FR4 module uses a dual CS connector to connect to 2 fiber pairs.

What does it mean when an electrical or optical channel is PAM-4 or NRZ?

NRZ stands for “Non Return to Zero” modulation, and describes an electrical or optical data channel where there are only two allowed amplitude levels (or symbols), with one amplitude level representing a digital ‘1’ and the other level representing a digital ‘0’. This is the predominate modulation scheme for transmitting data up to 25Gb/s, and is the simplest way to transmit digital data. The diagram below shows an example of an NRZ waveform, along with an ‘eye diagram’ for NRZ data. An eye diagram is simply a way of viewing a modulation scheme with each bit value superimposed on each other.



PAM-4 stands for Pulse Amplitude Modulation – 4, where ‘4’ refers to the number of different amplitude levels (or symbols) of the electrical or optical signal carrying the digital data. In this case, each amplitude level (or symbol) represents two bits of digital data. This enables a PAM-4 waveform to transmit twice as many bits as a NRZ waveform at the same symbol (or “toggle”) rate. The diagram below shows an example of a PAM-4 waveform, along with an eye diagram for PAM-4 data.



When a signal is referred to as “25Gb/s NRZ” or “25G NRZ”, it means the signal is carrying data at 25 Gbit / second with NRZ modulation. When a signal is referred to as “50G PAM-4”, or “100G PAM-4” it means the signal is carrying data at a rate of 50 Gbit / second, or 100 Gbit / second, respectively, using PAM-4 modulation.

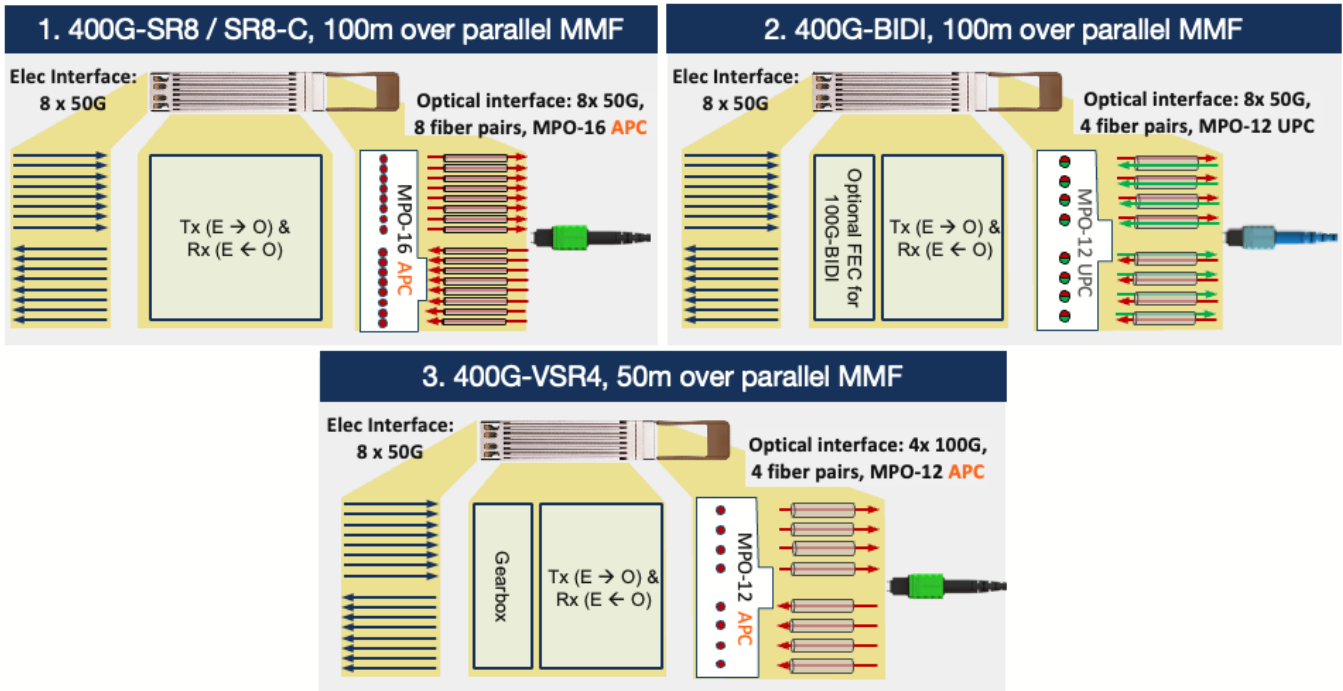
What is the reach, fiber type, connector, and optical modulation for each 400G transceiver type?

The table and block diagrams below summarize the key parameters for the different 400G transceivers supported in Arista’s platforms. Note that in all of the transceivers listed below, the electrical connector interface is always 8 x 50Gb/s PAM-4 (for a total of 400Gb/s). The optical signals however could be 8 x 50Gb/s PAM-4, or 4 x 100Gb/s PAM-4, depending on the type of module. For the modules that have 4 x 100Gb/s PAM-4 optical lanes, a gearbox chip inside the module converts the 8 x 50Gb/s PAM-4 electrical signals (from the board) to 4 x 100Gb/s PAM-4 signals required to modulate the optical signals.

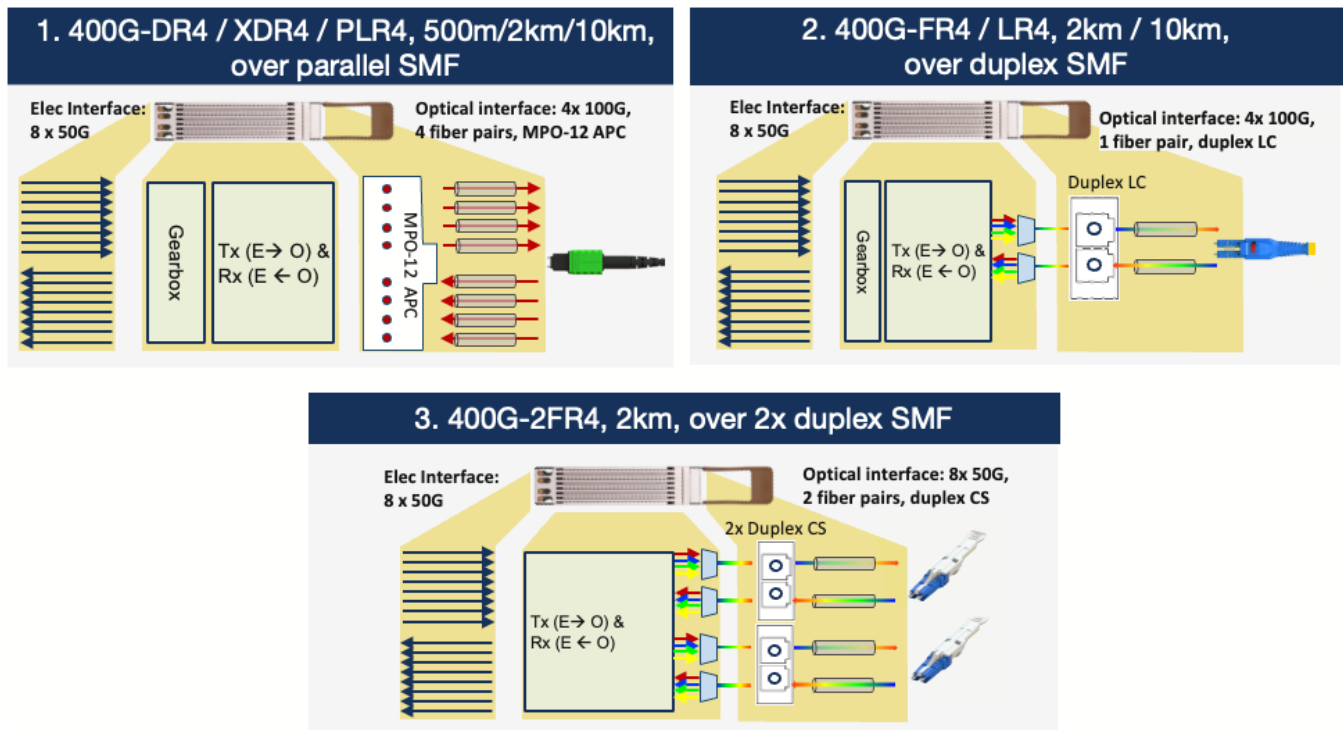
OSFP Part Number	QSFP-DD Part No.	Reach	Fiber type	Optical connector	No. of Fiber pairs	Optical waves / fiber	Optical modulation
OSFP-400G-SR8 & SR8-C	QDD-400G-SR8 & SR8-C	100m	Parallel MMF	MPO-16 Angled Connector (APC)	8	1	50G PAM-4
OSFP-400G-SRBD	QDD-400G-SRBD	100m	Parallel MMF	MPO-12 Un-angled Connector (UPC)	4	2	50G PAM-4
OSFP-400G-VSR4	QDD-400G-VSR4	50m	Parallel MMF	MPO-12 Angled Connector (APC)	4	1	100G PAM-4
OSFP-400G-DR4	QDD-400G-DR4	500m	Parallel SMF	MPO-12 Angled Connector (APC)	4	1	100G PAM-4
OSFP-400G-XDR4	QDD-400G-XDR4	2km	Parallel SMF	MPO-12 Angled Connector (APC)	4	1	100G PAM-4
OSFP-400G-PLR4	QDD-400G-PLR4	10km	Parallel SMF	MPO-12 Angled Connector (APC)	4	1	100G PAM-4
OSFP-400G-FR4	QDD-400G-FR4	2km	SMF	LC	1	4	100G PAM-4
OSFP-400G-LR4	QDD-400G-LR4	10km	SMF	LC	1	4	100G PAM-4
OSFP-400G-2FR4	-	2km	SMF	2 x CS	2	4	50G PAM-4

Block diagrams of each optic listed above are shown below.

400G Multimode Fiber Optics: Block Diagrams



400G Single mode Optics: Block Diagrams



What factors should be considered when choosing the right 400G transceiver for Multimode Fiber? What are the most common applications for each transceiver type?


Arista provides the following 400G optics for use in multimode fiber networks.

- i) **400G-SR8 / SR8-C** transceivers which use an MPO-16 APC optical connector and 8 pairs of MMF (16 total fibers), with a reach of 100m / 70m over OM4 / OM3 MMF.
- ii) **400G-BIDI** transceivers which use the widely deployed MPO-12 UPC optical connector and 4 pairs of MMF (8 total fibers) and, with a reach of 100m / 70m over OM4 / OM4 MMF.
- iii) **400G-VSR4** transceivers, which use an MPO-12 APC optical connector, and 4 pairs of MMF (8 total fibers), with a reach of 50m / 30m over OM4 / OM3 MMF.

The key features and common applications of each of these transceivers are described below.

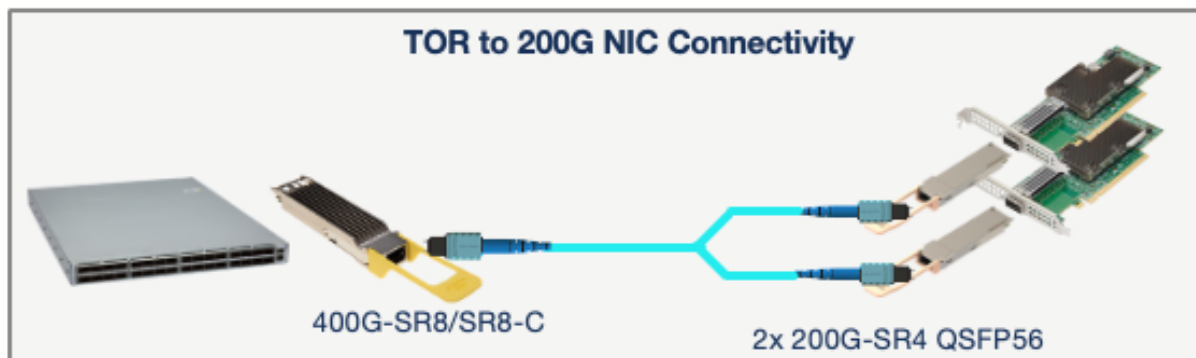
1. OSFP-400G-SR8 / SR8-C and QDD-400G-SR8 / SR8-C transceivers

The 400G-SR8 was the first 400G MMF transceiver available and has been deployed for point-to-point 400GE applications, such as leaf-to-spine connectivity, illustrated below.




While the 400G-SR8 provides cost-effective 400GE connectivity over MMF, it requires 16 fibers per transceiver and uses an MPO-16 APC fiber connector. Most 40G and 100G parallel MMF optics (such as the 40G-SR4 and 100G-SR4) use MPO-12 UPC fiber connectors. MPO-16 to 2x MPO-12 patch cables are required to use a 400G-SR8/SR8-C transceiver over an MPO-12 UPC based fiber plant – details of which can be found in the cabling guides available at: <https://www.arista.com/en/products/transceivers-cables/technical-resources>

Another key application for 400G-SR8 transceivers is optical breakout into 2x 200G-SR4 links, enabling TOR to host connectivity where 200G to the host is required, as illustrated below.

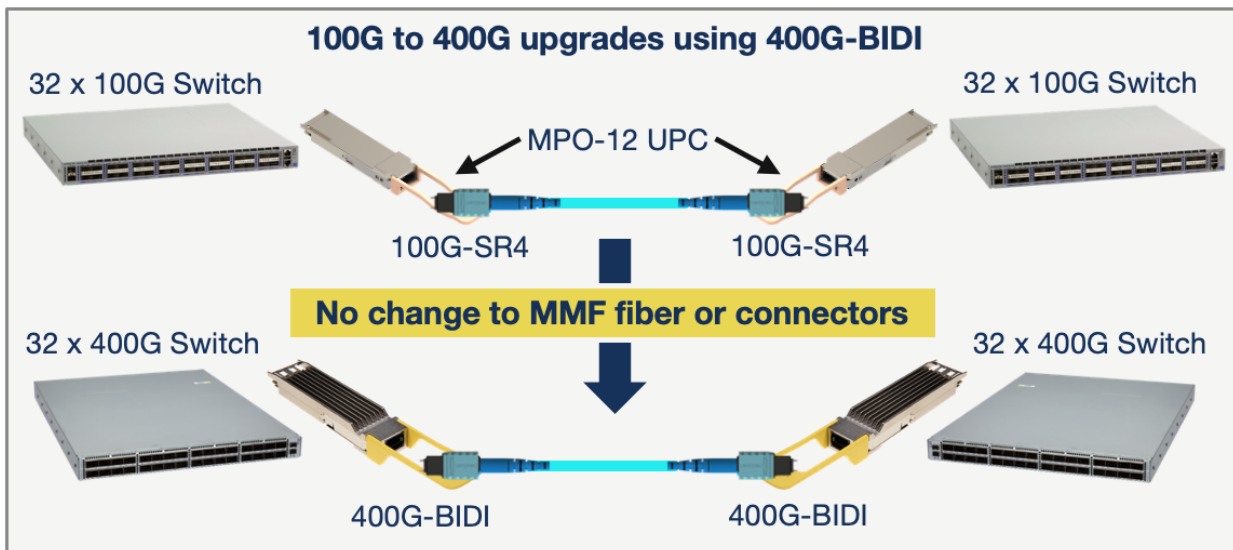


The 400G-SR8-C transceiver has the same features as the 400G-SR8, with the added ability to breakout into 8x 50G-SR or 8x 25G-SR optical links. It can therefore be used in applications that require high density 50G or 25G breakouts – as illustrated below.




2. OSFP-400G-SRBD and QDD-400G-SRBD, or “400G-BIDI” transceivers.

400G-BIDI transceivers use the widely deployed MPO-12 UPC connector for parallel multimode fiber. This allows existing 40G or 100G links that use 40G-SR4 or 100G-SR4 QSFP optics to be upgraded to 400GE with **no change to the fiber plant**, as illustrated below:



When configured for 400GE operation, the 400G-BIDI transceiver is compliant with the IEEE 400GBASE-SR4.2 specification for 400GE over 4 pairs of MMF.

Arista’s 400G-BIDI transceivers are also capable of breaking out into 4x 100GE links, and can be configured (via EOS) to interoperate either with the widely deployed base of 100G-BIDI (100G-SRBD) transceivers, or newer 100G-SR1.2 transceivers, as indicated below.



In summary, Arista's 400G-BIDI transceiver is software configurable to operate in any one of three operating modes:

- i) 400G-SR4.2 for point-to-point 400GE links
- ii) 4x 100G-BIDI for breakout and interop with 4x 100G-BIDI (100G-SRBD) transceivers
- iii) 4x 100G-SR1.2 for breakout and interop with 4x 100G-SR1.2 transceivers

More information on the 400G-BIDI transceiver is provided later in this document.

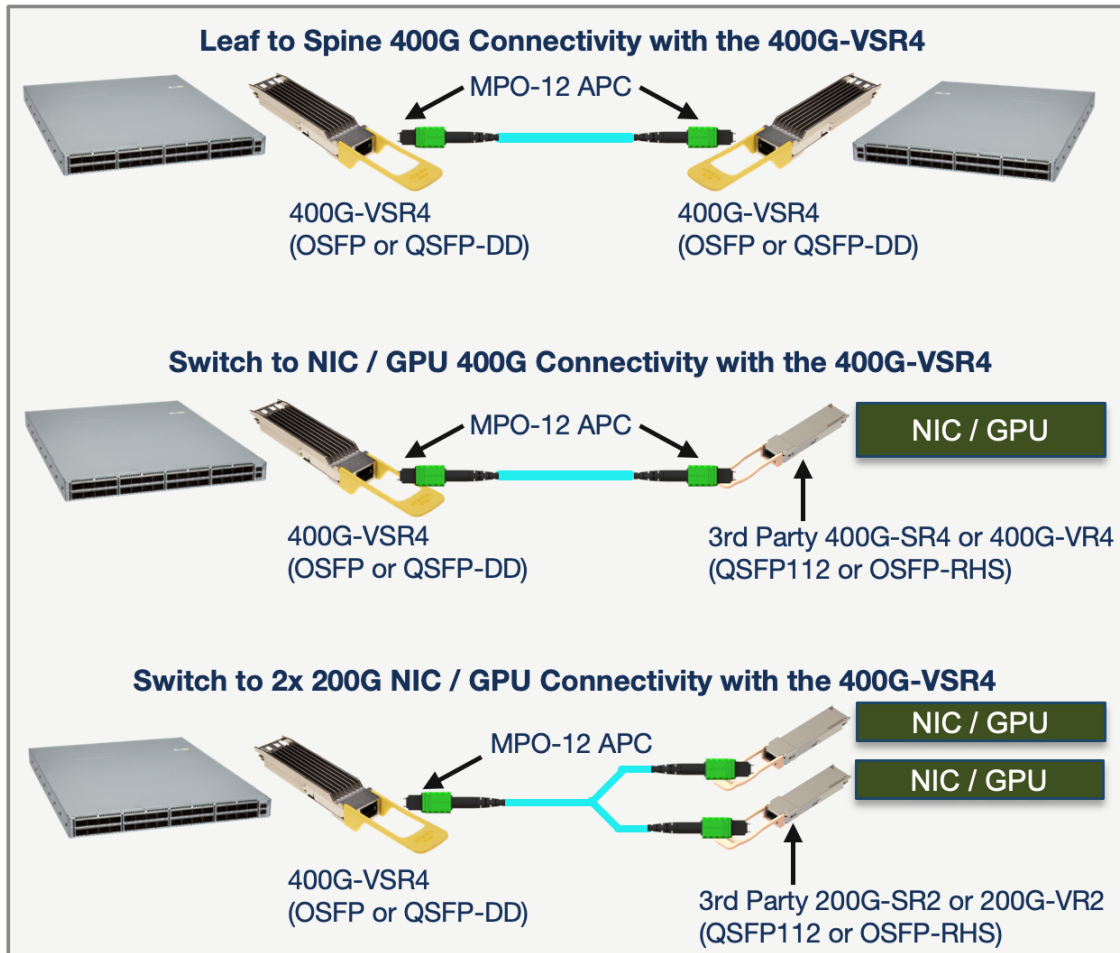
3. OSFP-400G-VSR4 and QDD-400G-VSR4 transceivers

400G-VSR4 transceivers are optimized for short reach (50m) links. Key applications of 400G-VSR4 modules include leaf-to-spine connectivity and switch-to-host connectivity when 400G connectivity to the host is required – such as AI clusters that use GPUs or NICs with 400G interfaces.

The QDD-400G-VSR4 and OSFP-400G-VSR4 modules use an 8x 50G PAM-4 electrical interface, and a 4x 100G PAM-4 optical interface (an internal gearbox does the 8x 50G to 4x 100G conversion). Many 400G GPU / NIC ports have a 4x 100G electrical interface and use form-factors such as a QSFP112 or an OSFP-RHS (an OSFP that uses a **Riding Heat Sink**, instead of an integrated heatsink). Arista's 400G-VSR4 modules will optically interop over 50m MMF with third-party QSFP112 or OSFP-RHS modules that are compliant to the IEEE 400GBASE-SR4 or 400GBASE-VR4 optical standards.

Some NICs that use QSFP112 or OSFP-RHS ports use only 2 out of the 4 possible lanes for 200G/port operation. Arista's 400G-VSR4 transceivers can break out into 2x 200G links and will interop over 50m MMF with QSFP112 or OSFP-RHS modules that are compliant to IEEE 200GBASE-SR2 or 200GBASE-VR2 optical standards.

Some typical applications for the 400G-VSR4 optics are summarized in the diagrams below.



400G-VSR4 transceivers use MPO-12 APC (Angled Physical Contact) fiber connectors to reduce optical back reflectance. Note that all 40G and 100G parallel MMF optics (such as the QSFP-40G-SR4 and QSFP-100G-SR4) use MPO-12 PC/UPC (un-angled) fiber connectors, so care must be taken to ensure that only MPO-12 APC connectors are used with OSFP-400G-VSR4 and QDD-400G-VSR4 modules.

What is the difference between 400G-BIDI, 400G-SRBD and 400G-SR4.2?

The terms 400G-BIDI, 400G-SRBD and 400G-SR4.2 refer to the same optical standard and are used interchangeably.

The IEEE 802.3 specification for 400G over 4 pairs of MMF using bidirectional technology is referred to as 400GBASE-SR4.2. In the term “400G-SR4.2”, the “SR” stands for “Short Reach” (100m), the “4” implies there are 4 fiber pairs used, and the “2” implies there are 2 wave per fiber (using bidirectional technology).

Because the 400G-SR4.2 utilizes Bi-Directional optical technology, it is also referred to as “400G-BIDI” or “400G-SRBD” (where the SRBD stands for Short Reach BIDI).

What FEC is required when the 400G-BIDI is configured for each of the three operating modes?

When the 400G-BIDI module is configured for either 400G-SR4.2 or 4x 100G-SR1.2 operation, FEC is implemented by the host silicon on the switch platform, and the 400G-BIDI module is put in “FEC bypass” mode. The FEC used for 400G-SR4.2 and 100G-SR1.2 is a Reed-Solomon (544,514) FEC, often referred to as KP-FEC.

When configured to interop with 4x 100G-BIDI (100G-SRBD) optics, Arista's 400G-BIDI modules are configured to terminate the host-to-module KP-FEC, and generate a separate 100G-BIDI FEC that will interop with the deployed base of 100G-BIDI (100G-SRBD) modules. The ability of Arista's 400G-BIDI modules to interop with 100G-BIDI (100G-SRBD) modules protects the substantial investments already made in 100G-BIDI deployments.

The 400G-BIDI module from Arista is capable of being broken out into 4x 100G-BIDI or 4x 100G-SR1.2 links. What is the difference between 100G-BIDI and 100G-SR1.2?

The difference between 100G-BIDI (100G-SRBD) and 100G-SR1.2 is the FEC (Forward Error Correction) used. 100G-BIDI (100G-SRBD) modules have been widely deployed for 100G operation over duplex MMF and use a FEC implementation that was developed prior to the IEEE standardization of KP-FEC for 50G PAM-4 based modules. Because of the differences in FEC implementation, 100G-SRBD and 100G-SR1.2 modules are not interoperable with each other.

Note that Arista's 400G-BIDI transceivers can be configured to breakout to either 4x 100G-BIDI (100G-SRBD), or 4x 100G-SR1.2 modules, providing maximum flexibility and investment protection.


What factors should be considered when choosing the right 400G transceiver for single mode fiber? What are the most common applications for 400G over SMF?

The most widely deployed 400G transceivers for SMF fall into two categories:

- i) 400G-FR4 / LR4 transceivers that use one pair of SMF with LC duplex connectors
- ii) 400G-DR4 / XDR4 / PLR4 transceivers that use 4 pairs of SMF and MPO-12 APC connectors.

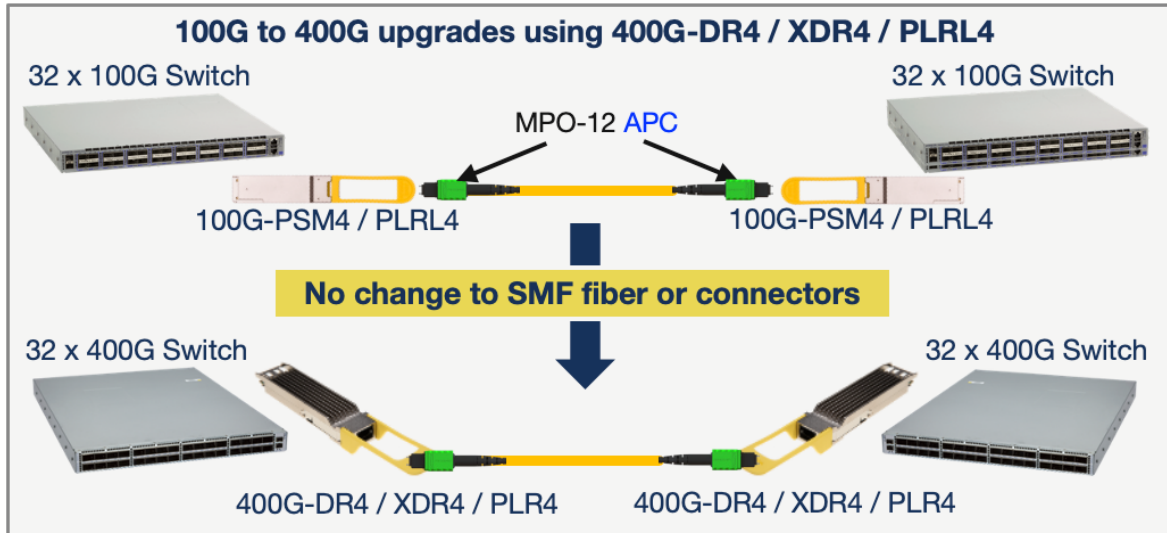
1. 400G-FR4 / LR4 transceivers

400G-FR4 and LR4 transceivers have reaches of 2km and 10km respectively and are widely deployed for point-to-point 400GE connectivity. 400G-FR4 / LR4 transceivers use the same fiber plant as all duplex 100G SMF optics (such as 100G-CWDM4, 100G-LR4 or 100G-DR/FR/LR optics), enabling easy 100G to 400G upgrades with no change to the fiber plant, as indicated below.

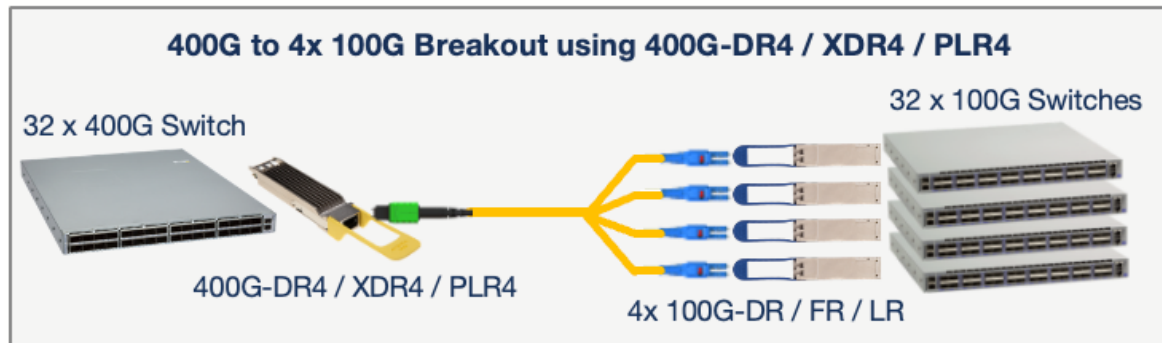


2. 400G-DR4 / XDR4 / PLR4 Transceivers

400G-DR4 / XDR4 / PLR4 transceivers can also be used for point-to-point 400GE connections and use the same fiber and connectors as 100G-PSM4 / PLRL4 transceivers, enabling 100G to 400G upgrades with no change to the fiber plant, as indicated below.



The 400G-DR4 / XDR4 / PLR4 transceivers are also capable of breaking out into 4x 100GE links, and are commonly used to breakout a single 400G port to 4x 100G-DR / FR / LR QSFPs, as shown below.



Can the 400G-DR4, 400G-XDR4 and 400G-PLR4 transceivers interoperate with each other?

Yes, the 400G-DR4 / XDR4 / PLR4 transceivers can interoperate with each other, with the max distance limited by the shortest reach transceiver.

Can the 400G-FR4 and 400G-LR4 transceivers interoperate with each other?

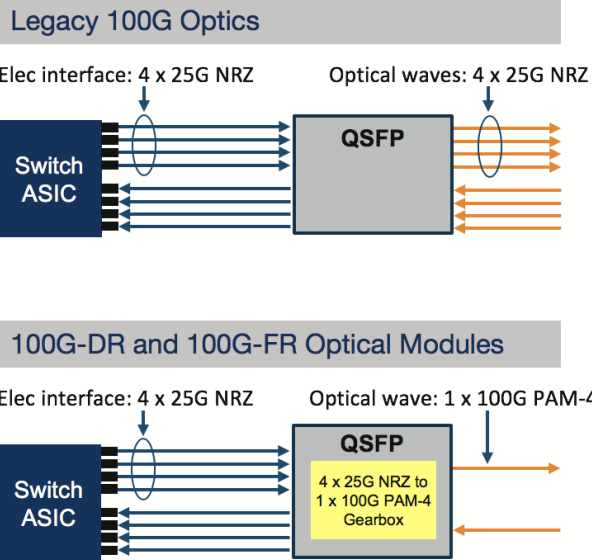
Yes, the 400G-FR4 and 400G-LR4 transceivers can interoperate up to a reach of 2km (limited by the FR4). Note that max allowed Receiver power for the 400G-FR4 (max Rx power of 3.5dBm) may require a minimum level of attenuation to be present if connected to a 400G-LR4 transmitter (Max Tx power of 5.1dBm).

What are the 100G-DR, 100G-FR and 100G-LR QSFP transceivers?

The 100G-DR / FR / LR transceivers are 100G QSFP modules that can be plugged into any Arista 100G QSFP port. The part numbers for these parts are **QSFP-100G-DR**, **QSFP-100G-FR** and **QSFP-100G-LR**.

The electrical connector interface is 4 x 25G NRZ, while the optical interface is a single 100Gbit/s PAM-4 optical signal. The 100G-DR / FR / LR modules use a gearbox to convert the 4 x 25G NRZ electrical signals to a 1 x 100G PAM-4 optical signal. This is in contrast to legacy QSFP100 modules (such as a CWDM4 or LR4 100G module), which have 4 x 25G NRZ optical wavelengths multiplexed onto one fiber.

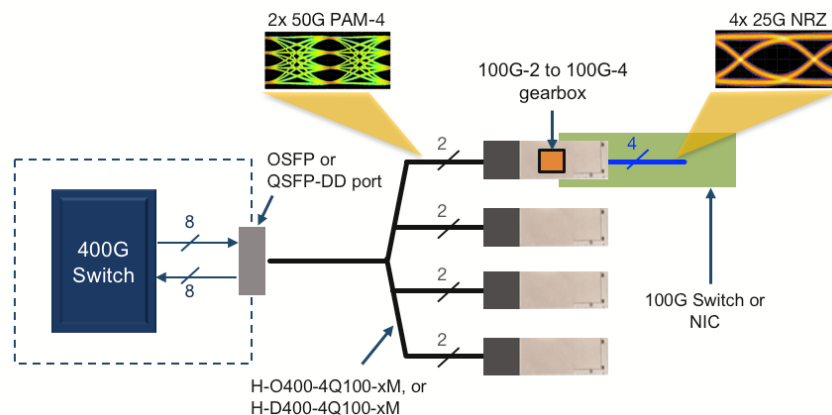
Because of the different optical modulation scheme, 100G-DR / FR / LR modules will not interoperate with legacy 100G modules (such as CWDM4, LR4 etc), but will interop with 400G-DR4 / XDR4 / PLR4 modules using breakout cables. The 100G-DR / FR / LR also interoperate with each other, up to the maximum distance of the shorter reach optics module (e.g. a 100G-DR will interoperate with a 100G-FR or LR for up to 500m in reach). The difference between legacy 100G QSFP module, and a 100G-DR / FR module is illustrated below.



What are Active Copper Breakout Cables?

Arista's active copper breakout cables (H-O400-4Q100-xM & H-D400-4Q100-xM) are active electrical cables that enable a 400G port to breakout into 4x 100G QSFP ports, over a length of up to 5m, where each 100G QSFP port uses 4x 25G NRZ electrical lanes (sometimes referred to as a 100G-4 interface).

The QSFP end of the active breakout cable includes a gearbox chip which converts 2x 50G PAM-4 electrical signals from the 400G port (a 100G-2 interface) into 4x 25G NRZ electrical signals (a 100G-4 interface). The Arista active copper breakout cables offers a cost-effective way to breakout a 400G port into 4x 100G and connect to an installed base of 100G QSFP ports.



What do the terms 100G-2, 100G-4, 50G-2, 200G-4 and 400G-8 mean?

These terms describe the bandwidth of an ethernet link, and the number of lanes used to achieve this bandwidth.

Every front panel port in an ethernet switch consists of one or more electrical lanes that transmit and receive ethernet data. For 10G SFP, 25G SFP or 50G SFP ports, a single electrical lane is used (in each direction) and modulated at 10G, 25G or 50G. For higher data rates, multiple lanes are required.

For example, a 100G QSFP port uses 4 lanes, with each lane running at 25Gb/s, or a '100G-4' interface. The number before the letter 'G' indicates the bandwidth of the ethernet link, and the number after the '-' indicates the number of data lanes required to achieve this bandwidth.

The table below summarizes the terminology used to describe common ethernet speeds, the number of lanes required to achieve this bandwidth, and some common applications of these interface types:

Ethernet link description	Link bandwidth	No. of lanes	Modulation of each lane	Common applications
10G	10Gb/s	1	10G NRZ	All 10G SFP+ ports
25G	25Gb/s	1	25G NRZ	All 25G SFP ports
50G	50Gb/s	1	50G PAM-4	All 50G SFP ports
50G-2	50Gb/s	2	25G NRZ	A 100G QSFP port on an Arista switch can often be configured to operate as 2x 50G-2 Ethernet links. 50G-2 QSFP interfaces are sometimes found in 50G NICs. Only 2 lanes (out of the 4 lanes available on a QSFP connector) are used. The Arista CAB-Q-2Q-100G-xM passive breakout cables can be used to breakout a 100G-4 QSFP port to 2x 50G-2 QSFP ports.
100G-4	100Gb/s	4	25G NRZ	All 100G QSFP ports
100G-2	100Gb/s	2	50G PAM-4	A 400G OSFP or QSFP-DD port on an Arista switch can be configured to operate as 4x 100G-2 Ethernet links. A 400G-DR4 / XDR4 / PLR4 optical module accepts 4x 100G-2 electrical links, and then uses a gearbox to convert 2x 50G PAM-4 electrical signals (100G-2) to / from a 1x 100G PAM-4 optical signal (100G-1). 100G-2 QSFP interfaces are sometimes used on 100G NICs. Only 2 electrical lanes (out of the 4 electrical lanes) are used. The Arista CAB-O-4Q-400G-xM and CAB-D-4Q-400G-xM passive breakout cables can be used to breakout a 400G OSFP or QSFP-DD port, respectively, to 4x 100G-2 QSFP ports.
100G-1	100Gb/s	1	100G PAM-4	The optical output of "100G single lambda" optical module is a single optical wavelength at a data rate of 100Gb/s using 100G PAM-4 modulation. A 100G-1 optical interface can be converted, using a gearbox, to a 100G-2 electrical interface (e.g. a 400G-DR4 optical module), or a 100G-4 electrical interface (e.g. a 100G-DR optical module)
200G-4	200Gb/s	4	50G PAM-4	All 200G QSFP ports. A 400G OSFP or QSFP-DD port on an Arista switch can often be configured to operate as 2x 200G-4 Ethernet links. The arista CAB-O-2Q-400G-xM and CAB-D-2Q-400G-xM passive breakout cables can be used to breakout a 400G OSFP or QSFP-DD port, respectively, to 2x 200G-4 QSFP ports.
400G-8	400G	8	50G PAM-4	All 400G OSFP and QSFP-DD ports

What does QSFP28, QSFP56 and SFP56 mean? What nomenclature should be used to describe the different types of QSFP and SFP ports?

The emergence of PAM-4 signaling has increased the types of interfaces available in QSFP and SFP form-factors. The table below summarizes how Arista describes each media type.

Arista Port / Module name	Other industry names	No. of Elec. lanes	Modulation	Comments
QSFP+ or 40G QSFP	QSFP+	4	10G NRZ	The QSFP was originally defined for <10G speeds. When it was adopted for 40G, it was called a QSFP+.
100G QSFP	QSFP28	4	25G NRZ	The "28" in QSFP28 refers to the max data rate (28Gb/s) each electrical lane can handle.
50G QSFP	50G QSFP28	2	25G NRZ	A 50G QSFP uses 2 out of the 4 available electrical lanes on a QSFP connector, with each lane running at 25Gb/s NRZ. It is used on some NICs and servers. No Arista products have 50G QSFP ports, but 100G QSFP ports can be connected to third party 50G QSFP ports using a copper breakout cable (DAC).
200G QSFP	QSFP56	4	50G PAM-4	A 200G QSFP uses 4 x 50Gb/s PAM-4 elec lanes. "56" refers to the the max data rate (56Gb/s) each electrical interface can handle.
50G SFP	SFP56	1	50G PAM-4	A 50G SFP that uses 1 x 50Gb/s PAM-4 elec lanes. This interface may be present on future NICs and servers.
25G SFP	SFP28	1	25G NRZ	The "28" refers to the maximum data rate on the single electrical lane, with 25G being the Ethernet rate
10G SFP or SFP+	SFP+	1	10G NRZ	SFPs were originally designed for < 10G speeds. When it was adopted for 10G, it became known as an SFP+.

How can I break out a 400G port and connect to 100G QSFP ports on existing platforms?

There are a few different ways to do this (as discussed earlier in this document), summarized below.


i) OSFP-400G-DR4 (or QDD-400G-DR4) to 4 x QSFP-100G-DR over 500m SMF.

Connect up to 4 x QSFP-100G-DRs to a single OSFP-400G-DR4 (or QDD-400G-DR4). The QSFP-100G-DR can plug into any Arista 100G QSFP port.




ii) OSFP-400G-XDR4 (or QDD-400G-XDR4) to 4 x QSFP-100G-FR over 2km SMF:

Connect up to 4 x QSFP-100G-FRs to a single OSFP-400G-XDR4 (or QDD-400G-XDR4). The QSFP-100G-FR can plug into any Arista 100G QSFP port.




iii) OSFP-400G-PLR4 (or QDD-400G-PLR4) to 4 x QSFP-100G-LR over 10km SMF

Connect up to 4x QSFP-100G-LRs to a single OSFP-400G-PLR4 (or QDD-400G-PLR4). The QSFP-100G-LR can plug into any Arista 100G QSFP port.



iv) H-O400-4Q100-xM (or H-D400-4Q100) to 4x QSFP100 ports with Active Copper DACs, 1m – 5m.

Connect up to 4x 100G QSFP ports to a single 400G OSFP or QSFP-DD port. The QSFP end of the active breakout DAC includes a gearbox chip which converts 2x50G PAM-4 electrical signals into a 4x 25G NRZ interface, the modulation format used in legacy 100G QSFP ports.



v) OSFP-400G-2FR4 to 2 x QSFP-100G-CWDM4 over 2km SMF:

If an OSFP port is configured for 2 x 100G (i.e. 200G total bandwidth), the OSFP-400G-2FR4 module can be used to connect to 2 x QSFP-100G-CWDM4 transceivers over duplex single mode fibers. Configuring an OSFP port for 200G total bandwidth means each of the 8 electrical lanes to/from the OSFP operate at 25Gb/s NRZ, the same modulation format used in legacy 100G QSFP ports.




vi) **OSFP-400G-SRBD (or QDD-400G-SRBD) to 4x QSFP-100G-SRBD or 4x 100G-SR1.2 QSFPs over 100m MMF**

Connect up to 4x QSFP-100G-SRBD, or 4x 100G-SR1.2 QSFPs to a single 400G-BIDI module.




vii) **OSFP-400G-SR8 (or QDD-400G-SR8) to 2 x QSFP-100G-SR4 QSFPs over 100m MMF:**

If an OSFP port is run at 200G total bandwidth, the OSFP-400G-SR8 module can be used to connect to 2 x QSFP-100G-SR4 transceivers using a multimode breakout cable.



viii) **Passive DAC breakout cable using CAB-O-2Q-400G-xM / CAB-O-2Q-200G-xM or CAB-D-2Q-400G-xM / CAB-D-2Q-200G-xM:**

If the OSFP or QSFP-DD port is run at 200G total bandwidth, a passive DAC breakout cable can be used to connect an OSFP or QSFP-DD port into 2x 100G QSFP ports.



What other breakout options are possible?

The tables below summarize the breakout options for 400G OSFP and QSFP-DD ports. In the tables below, the term “near end” transceiver refers to the transceiver / cable plugged into the OSFP or QSFP-DD port, and the “remote end” refers to the “broken out” transceiver / cable connected to the QSFP or SFP port. For completeness, these tables include breakouts to the 100G QSFP ports already described above.

Breakout options using OSFP or QSFP-DD optical modules.

Near end port	Near end port config	Fiber type	No. of fiber pairs	Optical Modulation	Remote end
OSFP(QDD)-400G-SRBD	4x 100G-2	Parallel MMF, 100m	4	50G PAM-4	4x QSFP-100G-SRBD or 4x 100G-SR1.2 QSFPs
OSFP(QDD)-400G-SR8 or OSFP(QDD)-400G-SR8-C	2 x 100G-4*	Parallel MMF, 100m	8	25G NRZ	2 x QSFP-100G-SR4
	2 x 200G-4	Parallel MMF, 100m	8	50G PAM-4	2 x QSFP-200G-SR4
OSFP(QDD)-400G-SR8-C	8x 50G	Parallel MMF, 100m	8	50G PAM-4	8x 50G-SR modules
	8x 25G*	Parallel MMF, 100m	8	25G NRZ	8x SFP-25G-SR modules
OSFP(QDD)-400G-VSR4	2x 200G-4	Parallel MMF, 50m	4	100G PAM-4	2x Third Party 200G-SR2 or 2x 200G-VR2 modules**
OSFP(QDD)-400G-DR4 / XDR4 / PLR4	4 x 100G-2	Parallel SMF, 500m / 2km / 10km	4	100G PAM-4	4x QSFP-100G-DR / FR / LR
OSFP-400G-2FR4	2 x 200G-4	2 x duplex SMF, 2km	2	50G PAM-4	2 x QSFP-200G-FR4
	2 x 100G-4*	2 x duplex SMF, 2km	2	25G NRZ	2 x QSFP-100G-CWDM4

* The OSFP / QSFP-DD port is configured for an aggregate bandwidth of 200G (instead of 400G).

** Third party 200G-SR2 or 200G-VR2 modules sometimes used in 200G QSFP112 or OSFP-RHS based NICs.

Breakout options using OSFP or QSFP-DD 400G copper DACs

Near end port	Near end port config	Electrical Modulation	Reach	Remote port
CAB-O-2Q-400G-xM or CAB-D-2Q-400G-xM	2 x 200G-4	50G PAM-4	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	2 x Third party 200G QSFP ports
	2 x 100G-4*	25G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	2 x 100G QSFP ports
	2 x 40G-4**	10G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	2 x 40G QSFP ports
CAB-O-4Q-400G-xM or CAB-D-4Q-400G-xM	4 x 100G-2	50G PAM-4	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	4 x 100G QSFP ports, using 2 x 50G PAM-4 lanes per QSFP
	4 x 50G-2	25G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	4 x Third party 50G QSFP ports, using 2x 25G NRZ lanes per QSFP
CAB-O-8S-400G-xM or CAB-D-8S-400G-xM	8 x 50G	50G PAM-4	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	8 x 50G SFP ports
	8 x 25G*	25G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	8 x 25G SFP ports
	8 x 10G**	10G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	8 x 10G SFP ports

* The OSFP / QSFP-DD port is configured for an aggregate bandwidth of 200G (instead of 400G).

** The OSFP / QSFP-DD port is configured for an aggregate bandwidth of 80G (instead of 400G).

Breakout options using OSFP or QSFP-DD 400G Active Copper DACs

Near end port	Near end port config	Electrical Modulation	Reach	Remote port
H-O400-4Q100-xM or H-D400-4Q100-xM	4 x 100G-2	50G PAM-4 at OSFP/QSFP-DD, to 25G NRZ at QSFP	1m, 2m, 3m, 5m	4 x 100G QSFP ports (with 4x 25G NRZ electrical interface)

Breakout options using OSFP or QSFP-DD 200G copper DACs

Near end port	Near end port config	Electrical Modulation	Reach	Remote port
CAB-O-2Q-200G-xM or CAB-D-2Q-200G-xM	2 x 100G*	25G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	2 x 100G QSFP ports
	2 x 40G**	10G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	2 x 40G QSFP ports
CAB-O-4Q-200G-xM or CAB-D-4Q-200G-xM	4 x 50G	25G NRZ	1m, 2m, 3m (OSFP) 1m, 3m, 2.5m (QSFP-DD)	4 x 50G QSFP ports, using 2x 25G NRZ lanes per QSFP
CAB-O-8S-200G-xM or CAB-D-8S-200G-xM	8 x 25G*	25G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	8 x 25G SFP ports
	8 x 10G**	10G NRZ	1m, 2m, 3m (OSFP) 1m, 2m, 2.5m (QSFP-DD)	8 x 10G SFP ports

* The OSFP / QSFP-DD port is configured for an aggregate bandwidth of 200G (instead of 400G).

** The OSFP / QSFP-DD port is configured for an aggregate bandwidth of 80G (instead of 400G).

What is the maximum supported distance for 400G Transceivers and Cables?

The maximum distance currently supported by Arista's 'client optics' 400G transceivers is 10km over SMF, using either the 400G-LR4 or 400G-PLR4 modules.

For passive copper DACs, a max length of 3m is supported. For active copper cables, a max of 5m is supported.

For DWDM links, Arista provides support for 400ZR coherent modules capable of closing 120km. The 400ZR modules integrate coherent DWDM transmission into the switch / router port, eliminating the requirement for expensive optical transport equipment. Please refer to Arista's [400ZR DWDM DCI Solution Datasheet](#)

What is the maximum power consumption of 400G OSFP and QSFP-DD transceivers?

The power consumption of 400G client transceivers will range from 7W to 12W per port. Coherent DWDM transceivers may draw up to as much as 24W per port. Refer to the transceiver data sheet for power consumption values of individual modules.

What industry standards are associated with each of the 400G Transceivers and Cables?

The table below summarizes the Arista 400G transceivers and cables and the associated industry standards.

Arista Products	Associated Industry Standard
CAB-O-O-400G-xM, and CAB-D-D-400G-xM	There is no official IEEE standard for 400G over copper, however the implementation is often referred to as 400G-CR8, and follows the electrical specifications of 200GBASE-CR4, as per IEEE 802.3cd
OSFP-400G-SRBD and QDD-400G-SRBD	400GBASE-SR4.2 as defined in IEEE 802.3-2022
OSFP-400G-SR8 / SR8-C and QDD-400G-SR8 / SR8-C	400GBASE-SR8, as defined in IEEE802.3cm. Breaks out to 8 x 50G-SR, as defined in IEEE 802.3cd
OSFP-400G-VSR4 and QDD-400G-VSR4	400GBASE-SR4 and 400GBASE-VR4, as defined in IEEE 802.3db. Arista's 400G-VSR4 transceivers are fully compliant with 400GBASE-VR4 and will optically interop with 400GBASE-SR4 up to a reach of 50m.
OSFP-400G-DR4, and QDD-400G-DR4	400GBASE-DR4, as defined in IEEE 802.3bs. Breakout to 100G-DR (IEEE 802.3cd)
OSFP-400G-XDR4, and QDD-400G-XDR4	Optical specs are identical to the 100G-FR, as defined in the 100G Lambda MSA. The 400G XDR4 essentially contains 4x 100G-FR links
OSFP-400G-PLR4, and QDD-400G-PLR4	Optical specs are identical to the 100G-LR, as defined in the 100G Lambda MSA. The 400G PLR4 essentially contains 4x 100G-LR links
OSFP-400G-FR4 / LR4, & QDD-400G-FR4 / LR4	400-FR4/LR4, as defined in the 100G Lambda MSA. http://100glambda.com/
OSFP-400G-2FR4	2 x 200GBASE-FR4, as defined in IEEE 802.3bs
QSFP-100G-DR	100GBASE-DR, as defined in IEEE 802.3cd
QSFP-100G-FR / LR	100G-FR / LR, as defined in the 100G Lambda MSA
All OSFP transceivers	The OSFP MSA: https://osfpmsa.org/
All QSFP-DD transceivers	The QSFP-DD MSA: http://www.qsfp-dd.com/

What are the pros and cons of using OSFPs or QSFP-DDs?

Arista platforms support both the OSFP and the QSFP-DD form factors – in other words, Arista will provide both OSFP and QSFP-DD compatible hardware. Both form factors allow for high density 36x 400G ports in 1RU.

The QSFP-DD form factor is based on the QSFP form factor with an extra row of electrical pins added. It allows strict backwards compatibility with 40G and 100G QSFP modules. To dissipate the larger power of 400G modules, the QSFP-DD relies on an external heatsink that is part of the switch platform – i.e. when a QSFP-DD module is inserted into a QSFP-DD port, the platform must provide a heatsink that makes contact with the QSFP-DD module with sufficient pressure to guarantee a low thermal resistance interface.

The OSFP form factor was designed from the ground up for optimal performance at 400G and 800G. It allows for backwards compatibility to 40G and 100G QSFP optics using the ADPT-O-Q-100G passive adapter. A key

difference of the OSFP is that the heatsink is integrated into the module case itself. This has the advantage of ensuring the best possible thermal contact between the power dissipating components and the heatsink, enabling better thermal performance. The OSFP also has approximately 50% more surface area than the QSFP-DD, enhancing the ability of the module to dissipate heat. OSFP modules run between 5C – 15C cooler than QSFP-DD modules in equivalent systems. This allows support for a large range of optics, and operating optics at lower temperature improves reliability.

What type of optical connectors do the 400G-FR4 / LR4, 400G-DR4 / XDR4 / PLR4, 400G-BIDI (400G-SRBD), 400G-SR8, 400G-VSR4 and 400G-2FR4 transceivers use?


The 400G-FR4/LR4 modules use duplex LC fiber connectors – the same connectors that are used on existing 100G and 40G QSFPs that use duplex SMF (e.g. 100G-CWDM4, 40G-LR4 etc).




The 400G-DR4 / XDR4 / PLR4 modules use a single row, angled (APC) MPO12 connector for use with parallel single mode fiber – which is the same connector used on 100G and 40G QSFPs that use parallel SMF (such as the 100G-PSM4, 40G-PLR4, etc).

The 400G-VSR4 modules use a single row, angled (APC) MPO12 connector for use with parallel multi mode fiber. Note that the 100G and 40G QSFPs that use parallel MMF (e.g. 100G-SR4 and 40G-SR4) use MPO12 UPC/PC (un-angled) connectors. Care must be taken to ensure that only MPO12 APC MMF connectors are used with 400G-VSR4 transceivers.


Although an MPO12 cable can have up to 12 fibers, only 8 out of the 12 fibers are used for the 400G-DR4 / XDR4 / PLR4 and 400G-VSR4 transceivers (4 for Tx and 4 for Rx). An image and drawing of an MPO12 connector are shown below:




The Arista 400G-XDR4-S and 400G-PLR4-S modules use “SN” connectors, which allow each of the four parallel individual links to be plugged and unplugged – eliminating the need for breakout cables.



400G-BIDI optical modules use a single row, un-angled (UPC) MPO12 multimode fiber connector. Although a MPO12 cable can have up to 12 SMF fibers, only 8 out of the 12 fibers are used. The un-angled (UPC) MPO-12 connector used by the 400G-BIDI module is identical to the MPO12 connectors used by the widely deployed 100G-SR4 modules – making for simple 100G to 400G migration using the 400-BIDI transceivers. An image, and drawing of the MPO-12 connector used for the 400G-BIDI modules is shown below:



The 400G-SR8 optics module uses a single row, angled (APC) MPO16 connector for use with parallel multi-mode fiber. All 16 fibers are used (8 Tx and 8 Rx). An image and drawing of an MPO16 connector are shown below:



The 400G-2FR4 optics module uses 2 x CS connectors, for use with single mode fiber. An image and drawing of 2 x CS connectors is shown below:



What does “APC” or PC/UPC mean when it comes to optical connectors? Which optical connectors use APC, and which use UPC?

PC and UPC refers to “Physical Contact” or “Ultra Physical Contact” fibers, while APC refers to “Angled Physical Contact” fibers. These terms refer to geometry of the end face of the optical fiber. In PC/UPC fiber connectors, the end face of the fiber is “flat”. With APC connectors, the end face of the fiber is polished at an angle to reduce back reflections.

The following table summarizes what type of fiber termination should be used with Arista optics.

Optics type (SMF or MMF)	Parallel or duplex fiber	Optical connector	Required fiber end polish	Comments
SMF and MMF	Duplex	Duplex LC	PC / UPC	All Arista optics that use duplex LC connectors should be connected to PC / UPC fibers.
SMF	Parallel	MPO-12	APC	All Arista optics with SMF MPO-12 connectors use APC fiber connections.
MMF	Parallel	MPO-12	UPC	All Arista 100G and 40G optics that use parallel MMF use MPO-12 UPC connectors. The Arista 400G-BIDI transceivers also use MPO-12 UPC connectors, enabling 100G-SR4 to 400G-BIDI upgrades without any change to fiber plant.
MMF	Parallel	MPO-12	APC	Arista 400G-VSR4 optics use MPO-12 APC connectors. Care must be taken to ensure that MPO-12 UPC connectors (used for 100G-SR4 and 400G-BIDI) are not used for 400G-VSR4 links.
MMF	Parallel	MPO-16	APC	The MPO-16 APC connector is used in the 400G-SR8 and 400G-SR8-C modules.

What additional resources are available on Transceivers and Cables?

Below is a list of additional resources available on the transceivers and cables page of www.arista.com.

Document	Description
Arista Transceivers and Cables Datasheet	Detailed specifications and ordering information
Transceiver and Cable Guide	Arista EOS support, physical attributes, laser safety and fiber cleaning instructions
400ZR DCI Solution Datasheet	Datasheet for Arista's 400G-ZR, 400G-ZR+ optical transceivers, and Arista's pluggable line system solutions.
100G Transceivers and Cables FAQ	100G Optics and cables frequently asked questions
40G Transceivers and Cables FAQ	40G Optics and cables frequently asked questions
Corning 400G Cabling Guide Leviton 100G/400G Cabling Guide Commscope 400G Cabling Guide	Partner documents such as fiber cabling guides from Cabling companies like Corning, CommScope, Leviton and Siemon