Fibre Types

Singlemode cables can be found in OS2 type. Multimode cables can be found in OM1, OM2, OM3, OM4 and OM5 types. Each type has different properties.

OS2 Singlemode Fibre

Unless requested otherwise all singlemode fibre cables they will conform to ITU-T G652.D. This standard describes the geometrical, mechanical and transmission attributes of the cable and ensures that the cable can be used in the operating wavelengths around 1310nm and 1550nm.

Also, in higher fibre density environments fibre with tighter cable routing and reduced bend radius there is another standard G.657.A1 and or G.657.A2. Cables meeting either of these standards are often referred to as Bend Insensitive (BI) or Reduce Bend Sensitive (RBS) fibre cable. All G.652.D, G.657.A1 and G.657.A2 all have the same physical size with internal and external core diameters of 9µm and 125µm, respectively, are completely compatible with each other and G.652 and G.657 are compatible with all standard connectors such as LC, SC, MU, and E2000 in UPC or APC polishes.

OM5 Multimode Fibre

OM5 New Standard developed when you are using Shortwave Wavelength Division Multiplexing (SWDM) applications. For 10 Gigabit and 100 Gigabit Ethernet it only transmits at the same distance as the low cost OM3 and OM4 Fibre. It has a suggested jacket colour of Lime Green. OM5 fibre, also known as WBMMF (wideband multimode fibre) at a minimum speed of 28Gbps per channel through the 850-953 nm window.

OM4 Multimode Fibre

It is a further improvement to OM3, uses a 50µm core and was developed specifically for VSCEL laser transmission and allows 10 Gigabit Ethernet up to 550m compared to 300M with OM3 and 100 Gigabit Ethernet at lengths up to 150 meters. Often supplied in a jacket colour of Erika Violet to differentiate it from OM3 . It is completely backwards compatible with OM3 fibre.

OM3 Multimode Fibre

Its core size is 50µm, but the cable is optimized for laser based equipment that uses fewer modes of light. As a result of this optimization, it is capable of running 10 Gigabit Ethernet at lengths up to 300 meters. Since its inception, production techniques have improved the overall capabilities of OM3 to enable its use with 40 Gigabit and 100 Gigabit Ethernet up to 100 meters. It has a suggested jacket colour of aqua.

Ethernet Cabling Guidelines

	Multimode Fibre OM3	Multimode Fibre OM4	Multimode Fibre OM5	Singlemode Fibre OS2 1310mm	Singlemode Fibre OS2 1550mm
10BASE-T (10Mb/s)	2000m	2000m		N/S	N/S
100BASE-T (100Mb/s)	2000m	2000m		N/S	N/S
1000BASE-T (GbE)	550m	1000m		5000m	N/S
10GBASE-T (10GbE)	300m	550m	400m	10000m	40000m
40GBASE-T (40GbE)	100m	100m	150m	10000m	40000m
100GBASE-T (100GbE)	100m	150m	150m	10000m	40000m

N/S = Not Specified

IEEE Standard	Application	Fibre Type	Baud Rate (GBd)	Transmission Type	# of λ's	# of Fibres	Distance (metres)	Connector
802.3cm	400GBASE-SR8	Multimode	50	Parallel	NA	16	70 (OM3) 100 (OM4)	MPO-16/MPO-24
802.3cm	400GBASE-SR4.2	Multimode	50	Parallel/SWDM	2	8	70 (OM3) 100 (OM4) 150 (OM5)	MPO-8/MPO-12
802.3bs	400GBASE-DR4	Singlemode	100	Parallel	NA	8	500	MPO-8/MPO-12
802.3bs	400GBASE-FR8	Singlemode	50	WDM	8	2	2,000	Duplex
802.3cu	400GBASE-FR4	Singlemode	100	WDM	4	2	2,000	Duplex
802.3cu	400GBASE-LR4-6	Singlemode	100	WDM	4	2	6,000	Duplex
802.3bs	400GBASE-LR8	Singlemode	50	WDM	8	2	10,000	Duplex

Current IEEE 400 Gb/s Ethernet standards for the data center using QSFP-DD transceivers.

Connectors

Fibre Optic Connectors

Unlike with copper cabling where one style of connector, the RJ45, dominates, with fibre there are a number of different connectors that can be used.

LC Connectors

The LC connector is sometimes referred to as a Small Form Factor (SFF) connector as it has much smaller dimensions than the SC connector. In fact it is possible to get twice as many fibres terminated on LC connectors in the space used by SC connectors. LC connectors have a simple latch, similar to that used on a RJ45 plug, making it familiar in use for network technicians. This has become the connector of choice in most new network fibre installations.

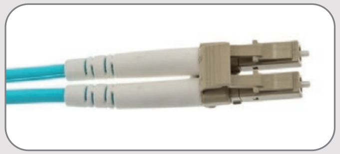

SC Connectors

The SC connector has a square format and has a push/pull latching mechanism. The image shows two SC connectors that have been joined using a special clip to create a SC Duplex connector with the A and B legs identified in the moulding of the clip.

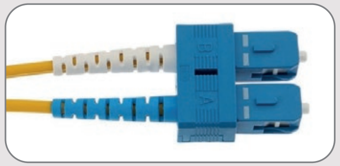

ST Connectors

The ST connector has a round barrel and uses a bayonet fixing mechanism to secure it to the mating connector. It is not recommended for new installations and is only seen in legacy fibre deployments.

MPO/MTP® Connectors

The MPO/MTP® connector has larger dimensions than the LC Duplex connector but can accommodate up to 24 fibres in a single ferrule, making it ideal for high density installation. This style of connector is also finding favour in multi-channel fibre applications such as 40G and 100G Ethernet. However, the high precision nature of the connector means that it is not suitable for field termination.

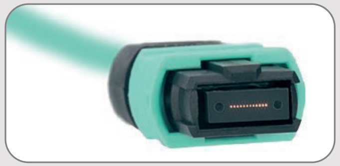

LC Duplex, CS, SN, Comparison

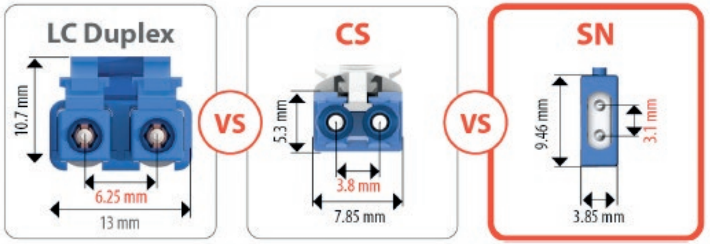

Comparison of traditional LC to new small form factor CS and SN duplex connectors.

Patch Cable Types Explained

Single Mode OS2 Patch Cable

OS2 Patch Cables are made to provide bandwidth used in long distance transmission. They are 2.8mm Duplex Design with white strain relief boots and our standard LSOH jacket colour is Yellow. Each fibre cable is tested for insertion loss. Individual test results are provided.

10 Gigabit 50µm OM3 and OM4 Patch Cable

OM3 Patch Cables are made to provide bandwidth that will support transmission above 10 Gigabits and can be used for applications up to 300 metres. They are 2.8mm Duplex Design with white strain relief boots and our standard LSOH jacket colour is Aqua Blue. Each fibre cable is tested for insertion loss. Individual test results are provided.

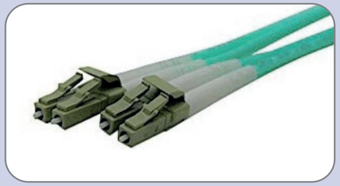

OM4 Patch Cables are made to provide bandwidth that will support transmission above 10 Gigabits and can be used for applications up to 550 metres. They are 2.8mm Duplex Design with white strain relief boots and our standard LSOH jacket colour is Aqua Blue or Erika Violet. Each fibre cable is tested for insertion loss. Individual test results are provided.

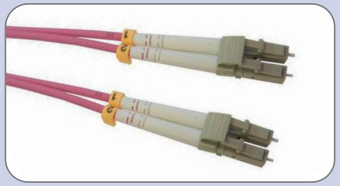

Uniboot LC Reverse Polarity Patch Cable

We can offer Uniboot LC Single and Multimode fibre patch cords offer a compact design and with a few simple steps the connector polarity can be reversed without any connector re-termination. Single mode OS2 has a Yellow LSOH jacket and Multimode OM3 and OM4 fibre with an Aqua Blue Jacket. Each fibre cable is tested for insertion loss. Individual test results are provided.

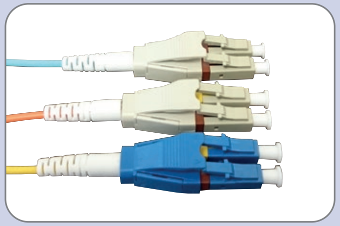

What's the Difference Between UPC Endface & APC Endface?

An APC connector endface features an 8-degree angle to direct reflected light into the cladding where it does not propagate back to the transmitter, improving system performance for 400 Gb/s PAM4 applications.

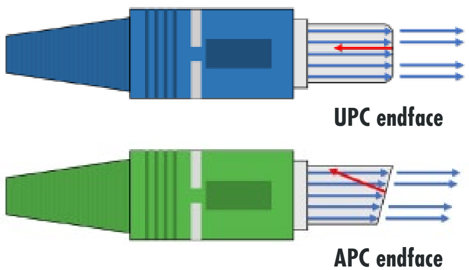

MPO/MTP

MPO/MTP (multi-fibre termination push-on) terminated cables are widely used in high-density data centres. The terms MPO and MTP are often used; however, MPO is the recognised fibre connector type, while MTP is a registered trademark of an MPO connector manufactured by US Conec. All MTPs are MPOs, but not all MPOs are MTPs. To the naked eye, there is little difference between the two connectors, and in cabling and transceiver interfaces, they are compatible with each other.

Note: MTP® is a registered trademark of US Conec Ltd. This is the term US Conec uses to describe their connectors. American Conec MTP products are fully compliant with MPO standards. Therefore, the MTP connector is a type of MPO connector.

MPO/MTP Design

MPO/MTP connectors are typically available with 8, 12, or 24 fibres, with 32, 48, 60, or 72 fibres available. However, these higher fibre counts are generally reserved for speciality super high-density multi-fibre arrays in large-scale optical switches.

MPO/MTP connectors are either male (with alignment pins) or female (without alignment pins). To join two MPO/MTP connectors together through an adapter, one connector must have alignment pins, and the other must not. The role of the alignment pins is to ensure that the fibres align perfectly, ensuring successful mating.

MPO/MTP connectors are often colour-coded to help contractors distinguish between the different fibre types and polish specifications for single mode. MPO/MTP connectors are made for both single-mode and multi-mode multi-fibre cables. The cable jacket is yellow in single-mode (OS2 applications), per Telecommunications Industry Association's (TIA) specification. The connector colour will vary depending on the connector type. Ultra-Physical Contact (UPC) connectors will also be yellow, while Angled Physical Contact (APC) connectors are green. With multi-mode OM3/OM4, both MPO/MTP connectors and cable jackets will be aqua, per TIA specifications.

MPO/MTP Connectors

Until recently, the primary method of connecting switches and servers within the data centre involved cabling organised around 12 or 24 fibres, typically using MPO/MTP connectors. The introduction of octal technology (eight switch lanes per switch port) enables data centres to match the increased number of ASIC I/Os (currently 256 per switch ASIC) with optical ports. Fibre cable manufacturers quickly reacted to this new dilemma and started to develop "Base-8" solutions. These new systems migrated the base unit of fibre bundles from 12 fibres to 8 fibres, and MPO/MTP cabling systems started to emerge based on 8 fibre MPO/MTP and trunk cables in multiples of 8 fibres.

New Base-8 and 16-fibre configurations support eight server connections using eight fibres to transmit and eight fibres to receive. MPO/MTP 16-fibre connectors are keyed differently to prevent connection with the 12-fibre MPO/MTP connectors.

The next speeds have now emerged as 400G. There will likely be minimal adoption of multi-mode fibre cabling, which will utilise an SR8 solution – 8 lanes of 50Gb/s and require a 16 fibre MPO/MTP connector. Base-8 infrastructure can support this 16-fibre solution by combining two Base-8 MPO/MTP connectors per link.

From a connectivity perspective, existing 8-fiber and 12-fiber MPO/MTP (Base-8/12) connectors will support 400 Gb/s 8-fibre applications, including 400GBASE-SR4.2 and 400GBASE-DR4. However, 16-fibre applications like 400GBASE-SR8 are now driving the need for 16-fiber MPO/MTP (Base-16) connectors, which have a different form factor from the Base-8/12 with an offset key.

MPO/MTP CONNECTORS IN THE CHANNEL

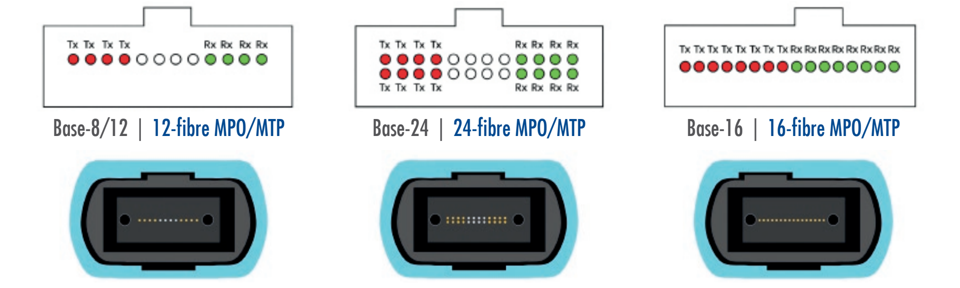

MPO/MTP Polarity

Achieving polarity is another challenge with MPO/MTPs. Polarity defines the direction of the light path or flow and is called the A-B-Cs of fibre polarity. This is more complex with multi-fibre MPO/MTP cables and connectors. Industry-standard TIA-568.3-D names three different polarity methods for MPO/MTPs: Method A, Method B, and Method C. Each method uses different types of MPO/MTP cables.

When examining a typical 12 fibre configuration, Method A uses a key-up connector on one end and a key-down connector on the other end so that the fibre located in Position 1 arrives at Position 1 at the other end. Method B uses key-up connectors on both ends to achieve the transceiver-receiver flip so that the fibre located in Position 1 arrives at Position 12 at the opposite end, the fibre located in Position 2 arrives at Position 11 at the opposite end, and so on.

Method C uses a key up connector on one end and a key down on the other end like Method A, but the flip happens within the cable itself where each pair of fibres is flipped so that the fibre in Position 1 arrives at Position 2 at the opposite end, and the fibre in Position 2 arrives at Position 1.

Whichever method is chosen, polarity needs to be consistent.

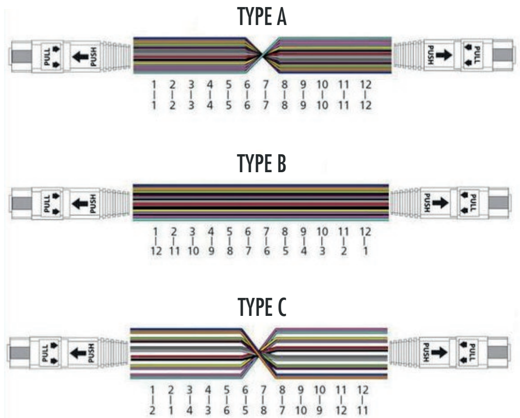

What Is Base-8 and 16 MPO/MTP Cabling?

As we forward through the evolution of Ethernet speeds for enterprise applications, with the migration of Ethernet speeds from 10Gb/s to 40Gb/, the IEEE802.3 standard launched the next Ethernet speed of 40Gb/s as a parallel transmission of four lanes of 10Gb/s (40GBASE-SR4), meaning that a single 40Gb/s fibre port on a switch would transmit and receive four separate lanes of 10Gb/s across 4 pairs of fibres. The standard transmission protocol had shifted from a legacy 2-fibre duplex system (one transmit and one receive) to an 8-fibre system (four transmits and four receives), and Base-8 was born.

These new 40Gb/s transceivers, called QSFP+ (Quad Small Form-factor Pluggable), had an MPO/MTP as the connector interface, utilising a standard Base-12 fibre connector but leaving the four middle connectors dark. Fibre cable manufacturers quickly reacted to this new dilemma and developed "Base-8" solutions. These new systems migrated the base unit of fibre bundles from 12 fibres to 8 fibres, and MPO/MTP cabling systems started to emerge based on 8-fibre MPO/MTP and trunk cables in multiples of 8 fibres.

The 100GBASE-SR4 standard utilised 4 lanes of 25Gb/s across 4 pairs of fibres, a further migration to 100Gb/s using Base-8 fibre cables and further strengthening the case for Base-8 connections.

The next speeds have now emerged as 400G. There will likely be minimal adoption on multi-mode fibre cabling, utilising an SR8 solution – 8 lanes of 50Gb/s and requiring a 16 fibre MPO/MTP connector. Base-8 infrastructure can support this 16-fibre solution by combining two Base-8 MPO/MTP connectors per link. The vast majority of 400G systems will run over single-mode fibre with either multiple wavelengths across a duplex fibre connection or four lanes of 100G transmitted across four pairs of fibres – again utilising Base-8 cabling.

From a connectivity perspective, existing 8-fiber and 12-fiber MPO/MTP (Base-8/12) connectors will support 400 Gb/s 8-fibre applications, including 400GBASE-SR4.2 and 400GBASE-DR4. However, 16-fibre applications like 400GBASESR8 are now driving the need for 16-fiber MPO/MTP (Base-16) connectors, which have a different form factor from the Base-8/12 with an offset key. While 24-fibre MPO/MTP connectors can be used for 16-fibre applications, their larger form factor and eight unused fibres make the Base-16 a more cost-effective solution.

While 100 GBd rate will enable 400 Gb/s over eight fibres supported via Base-8/12 connectors, future 800 Gb/s applications will sustain the need for Base-16 connectors. The IEC is currently working on specifications for Base-16 connectivity, and the IEEE Beyond 400 Gb/s Ethernet Study Group is looking to standardise on the Base-16 for future 800GBASE-SR8 and 800GBASE-DR8 applications.

The variety of connector technology options provides more ways to break out and distribute the additional capacity that octal modules offer. Connectors include parallel 8-, 12-, 16- and 24-fiber multi-push on (MPO/MTP) and duplex fibre LC, SN, MDC and CS connectors.

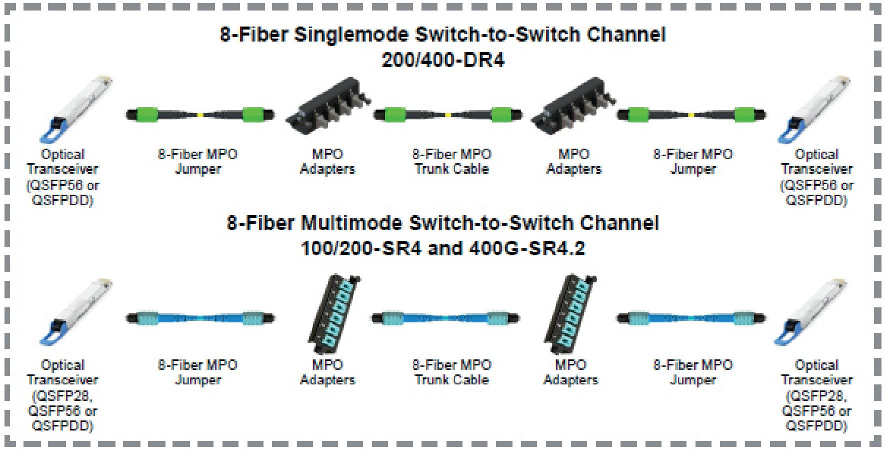

What Is Base-12 and 24 MPO/MTP Cabling?

Since manufacturers first produced fibre-optic cables, they came in multiples of twelve fibres. Cable fibre counts were commonly 12, 24, 48, 96, and 144, leading to systems around fibre optic cables also based on 12 fibres. For example, patch panels are typically based on multiples of 12, 24, 48, and 144, and network switches are generally based on multiples of 24.

The first MPO/MTP connector was with 12 fibres, and when in 1996, IBM introduced the first fibre cabling system based on 36, 72 and 144 fibre trunks terminated with 12 fibre MPO/MTP connectors. This system allowed the flexibility of being able to connect different "harness" types of assemblies to an MPO/MTP connector which fanned out to the duplex connection required by the equipment at the end of the cable, securely set on the magic number of 12!

As we forward through the evolution of Ethernet speeds for enterprise applications, with the migration of Ethernet speeds from 10Gb/s to 40Gb/s, there is a fundamental shift in transmission modes and the basis of fibre systems. The IEEE802.3 standard launched the next Ethernet speed of 40Gb/s as a parallel transmission of four lanes of 10Gb/s (40GBASE-SR4), meaning that a single 40Gb/s fibre port on a switch would transmit and receive four distinct lanes of 10Gb/s across 4 pairs of fibres. The standard transmission protocol had shifted from a legacy 2 fibre duplex system (one transmit and one receive) to an 8-fibre system (four transmit and four receive), and Base-8 was born.

A Base-12 MPO/MTP cabling uses fibre optical links based on increments of 12 fibres and 12-fibre MPO/MTP fibre optic connectors. Base-12 and or Base-24 are mainstream in MPO/MTP cabling systems.

A Base-24 MPO/MTP trunk can be used in a channel with both parallel transceivers at the uplink and duplex transceivers at the downlink. This configuration could be used to breakout 100 Gb/s to 25 Gb/s or 400 Gb/s to 100Gb/s in the future, without making any changes to the cabling system. The graphic below shows one channel breakout, three times the number of links that can be supported with the same cassettes and Base-24 trunk when fully populated. Alternate configurations are possible with other cabling schemes. Alternate configurations are possible with Base-8 or Base-12, using additional trunk cables to support an equivalent number of connections.

8-Fibre Channel Breakout with a Base-24 MPO Trunk

Fibre Utilisation

Base-12 is more advantageous in high-density network cabling because it has more cores and higher density. However, Base-8 MPO/MTP cabling has better fibre utilisation than Base-12 MPO/MTP cabling. Although the Base-12 MPO/MTP cabling is still the most common choice for most data centre operators, there are still no standardised transceivers using all 12 fibres. Many popular transceivers (SR4 transceivers, for example), only use an 8-fibre interface. Base-8 MPO/MTP links allow customers to directly connect fibres to those transceivers without any fibre waste. If we use a 12-fibre connector into a transceiver which only requires eight fibres, which means that four fibres are unused. You may consider using a conversion cable to convert the Base-12 cabling to Base-8 cabling (two Base-12 to three Base-8, for example) to make use of all fibres. However, it will cause additional insertion loss, reducing the cable performance.

In addition, when using MPO/MTP to LC duplex breakout harnesses to connect to switch line cards, the Base-8 harnesses easily route to all common port count line cards, as all common line cards contain a number of ports wholly divisible by the number four (since a Base-8 harness provides four LC duplex connections). If you use Base-12 MPO/MTP to LC duplex breakout harnesses, you will get 6 LC duplex connections. But these harness cables can't fully connect to line cards with 16 or 32 ports since 16 and 32 are not wholly divisible by 6.

Therefore, for the sake of fibre utilisation, Base-8 MPO/MTP cabling is a better choice. Of course, if you don't mind the fibre wasting, Base-12 MPO/MTP cabling could also be selected.

For Future Cabling System

Base-12 and Base-8 MPO/MTP cabling can be seamlessly converted to Base-2 cabling in relatively smaller networks like 10G connections. Under such circumstances, either is feasible. However, for more extensive networks like 40G, 100G or even 400G connection, a Base-8 solution will gain more widespread market acceptance since a more significant number of 40G, and 100G circuits are deployed utilising eight-fibre transceivers. So do 400G direct connections, which use 400G QSFP-DD transceivers and MPO/MTP-16 trunk cables. In addition, the 8-fibre cabling could work with the 24-fibre cabling seamlessly since a single 24-fibre MPO/MTP could break out to three 8-fibre MPO/MTPs. Customers deploying 10G data rates today, can still deploy the Base-8 system since upgrading to a 40G or 100G circuit will be much simpler and more cost-effective. Base-12 connectivity is not optimal for 8- fibre transceiver systems.

Current and Upcoming IEEE Ethernet Standards for 25 to 400 GB/s Targeted for the Enterprise Data Centre Environment

IEEE Standard	Transceiver	Signaling	Lane Rate	# of Lanes	# of Fibres	Fibre Type	Technology	Distance in metres
802.3bm	100G-SR4	NRZ	25 GB/s	4	8	Multimode	Parallel Optics	70 (OM3) 100 (OM4)
802.3cd	100G-SR2	PAM4	50 GB/s	2	4	Multimode	Parallel Optics	70 (OM3) 100 (OM4)
802.3cd	100G-DR	PAM4	100 GB/s	1	2	Singlemode	Short-Reach Duplex	500
802.3db*	100G-SR	PAM4	100 GB/s	1	2	Multimode	Duplex	70 (OM3) 100 (OM4)
802.3db*	100G-VR	PAM4	100 GB/s	1	2	Multimode	Short-Reach Duplex	30 (OM3) 50 (OM4)
802.3cd	200G-SR4	PAM4	50 GB/s	4	8	Multimode	Parallel Optics	70 (OM3) 100 (OM4)
802.3bs	200G-DR4	PAM4	50 GB/s	4	8	Singlemode	Short-Reach Parallel Optics	500
802.3db*	200G-SR2	PAM4	100 GB/s	2	4	Multimode	Parallel Optics	70 (OM3) 100 (OM4)
802.3db*	200G-VR2	PAM4	100 GB/s	2	4	Multimode	Short-Reach Parallel Optics	30 (OM3) 50 (OM4)
802.3cm	400G-SR8	PAM4	50 GB/s	8	16	Multimode	Parallel Optics	70 (OM3) 100 (OM4)
802.3cm	400G-SR4.2	PAM4	50 GB/s	8	8	Multimode	SWDM & Parallel Optics	70 (OM3) 100 (OM4) 150 (OM5)**
802.3bs	400G-DR4	PAM4	100 GB/s	4	8	Singlemode	Short-Reach Parallel Optics	500
802.3db*	400G-SR4	PAM4	100 GB/s	4	8	Multimode	Parallel Optics	60 (OM3) 100 (OM4)
802.3db*	400G-VR4	PAM4	100 GB/s	4	8	Multimode	Short-Reach Parallel Optics	30 (OM3) 50 (OM4)

** Currently in development.*
*** OM5 multimode fibre offers specified bandwidth at the 953-nanometer wavelength for SWDM applications. OM4 and OM5 multimode fibre offers the same bandwidth at the 850-nanometer wavelength and therefore support the same distances for all applications except for the SWDM 400G-SR4.2 application.*

QSFP Form Factor Transceivers

To support higher signalling rates, pluggable optical transceivers have also evolved. Iterations of QSFP transceiver modules have advanced from QSFP+ at a 10 GBd rate, to QSFP28 at a 25 GBd rate, and to QSFP56 at a 50 GBd rate. Equipped with a four-lane electrical interface, these transceivers enable 40, 100, and 200 Gb/s transmission respectively. The latest QSFP-DD form factor features an eight-lane interface, hence the DD designation for double density over existing four-lane QSFP form factors.

With the ability to support a 50 GBd rate, QSFP-DD transceivers support up to 400 Gb/s transmission, while enabling backwards compatibility with previous QSFP transceiver modules to avoid full active equipment replacement. Another eight-lane pluggable transceiver is the OSFP that also supports a 50 GBd rate but is not compatible with the QSFP form factor. The development of 100 GBd PAM4 is underway within the IEEE industry standards association to support higher speeds, but a 100 GBd rate has already been achieved by leveraging QSFP-DD transceivers and converting eight lanes of 50 Gb/s to four lanes of 100 Gb/s.

Just as the adoption of the QSFP28 form factor drove the adoption of 100G by offering high density and lower power consumption, the jump to 400G and 800G is being enabled by new transceiver form factors. The current SFP, SFP+ or QSFP+ optics are sufficient to enable 200G link speeds. However, making the jump to 400G will require doubling the density of the transceivers.

QSFP-Double Density (QSFP-DD7) and octal (2 times a quad) small form factor pluggable (OSFP8) Multi Source Agreements (MSAs) enable networks to double the number of electrical I/O connections to the ASIC. This not only allows summing more I/ Os to reach higher aggregate speeds, it also allows the total number of ASIC I/O connections to reach the network. The 1U switch form factor with 32 QSFP-DD ports matches 256 (32x8) ASIC I/Os. In this way, we can build high-speed links between switches (8*100 or 800G) but also have the ability to maintain the maximum number of connections.

The optical market for 400G is being driven by cost and performance as OEMs try to dial into the sweet spot of hyperscale and cloud scale data centers. In 2017, CFP8 became the first-generation 400G module form factor to be used in core routers and DWDM transport client interfaces. The CFP8 transceiver was the 400G form factor type specified by the CFP MSA. The module dimensions are slightly smaller than CFP2, while the optics support either CDAUI-16 (16x25G NRZ) or CDAUI-8 (8x50G PAM4) electrical I/O. As for bandwidth density, it respectively supports eight times and four times the bandwidth density of CFP and CFP2 transceiver.

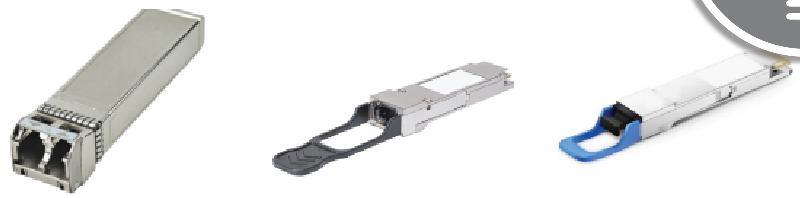

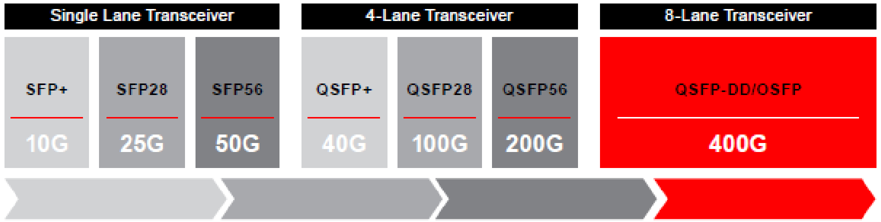

Transceiver technology has advanced to support 50 Gb/s over a single lane, 200 Gb/s over four lanes, and 400 Gb/s over eight lanes.

Current Versus Future Network Configurations

	Enterprise Data Centers		Cloud Data Centers	
	SERVER	UPLINKS	SERVER	UPLINKS
Current Network Speeds	1G	10G	10G	40G
Future Network Speed Options	10G OR 25G	40G OR 100G	25G OR 50G OR 100G	100G OR 200G OR 400G

Predictions – OSFP vs QSFP-DD

With regards to OSFP versus QSFP-DD, it's too early to tell which way the industry will go right now, both form factors are supported by leading data center Ethernet switch vendors and both have large customer support. Perhaps the enterprise will prefer QSFP-DD as an enhancement to current QSFP based optics. OSFP seems to be pushing the horizon with the introduction of OSFP-XD, extending the number of lanes to 16 with an eye toward 200G lane rates in the future. For speeds up to 100G, QSFP has become a go-to solution because of its size, power and cost advantage compared to duplex transceivers. QSFP-DD builds on this success and provides backwards compatibility which allows the use of QSFP transceivers in a switch with the new DD interface. Looking to the future, many believe that the 100G QSFP-DD footprint will be popular for years to come. OSFP technology may be favoured for DCI optical links or those specifically requiring higher power and more optical I/Os.

SN & MDC Connectors

Four SN and MDC small form factor duplex connectors fit into a single QSFP-DD transceiver to support 4X100 Gb/s breakout application in leaf-spine switch-to-switch links.

Note: MDC® is designed and manufactured by US Conec Ltd. The SN is designed and manufactured by Senko Advance Co Ltd (Japan).

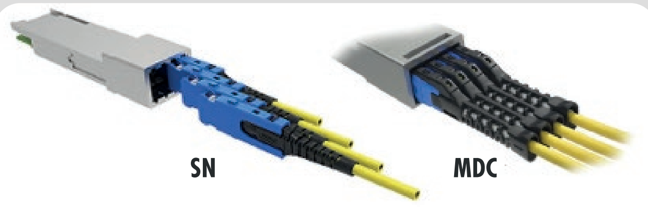

DATACENTRE PRODUCTS | TECH TERMS