

# The migration to 400G/800G: the Fact File



## Contents

<b>Planning to meet future data center challenges starts today. The Ethernet roadmap</b> .....	<b>3</b>
The DC landscape is evolving .....	3
<b>Timing the move to 400G/800G</b> .....	<b>4</b>
What does it mean for data centers? .....	4
Who (and what) is driving the change to higher speeds? .....	5
The transition to 400G is a group effort .....	5
<b>The Four Pillars of 400G/800G migration</b> .....	<b>6</b>
Higher switch-port densities (radix) – The end of ToR? .....	6
Transceiver technologies .....	7
<b>New transceiver formats</b> .....	<b>8</b>
<b>Modulation schemes</b> .....	<b>8</b>
<b>Predictions – OSFP vs QSFP-DD</b> .....	<b>9</b>
MPO and duplex connectors .....	9
<b>MPO connectors</b> .....	<b>10</b>
<b>Duplex connectors</b> .....	<b>12</b>
<b>Connector speed limits?</b> .....	<b>12</b>
Cabling advances .....	12
<b>Further information</b> .....	<b>13</b>
Additional helpful DC resources .....	13
<b>Additional helpful DC resources</b> .....	<b>13</b>
Why choose CommScope? .....	14
Products for your migration to 400G/800G .....	15

# Planning to meet future data center challenges starts today. The Ethernet roadmap

Across data centers the ground is shifting—again.

Accelerating adoption of cloud infrastructure and services is driving the need for more bandwidth, faster speeds and lower latency performance. Advancing switch and server technology are forcing changes in cabling and architectures. Regardless of your facility's market or focus, you need to consider the changes in your enterprise or cloud architecture that will likely be necessary to support the new requirements. That means understanding the trends driving the adoption of cloud infrastructure and services, as well as the emerging infrastructure technologies that will enable your organization to address the new requirements. Here are a few things to think about as you plan for the future.

Figure 1: Ethernet roadmap



## The DC landscape is evolving

### Global data usage

Of course, at the heart of the changes are the global trends that are reshaping consumer expectations and demand for more and faster communications, such as:

- Explosive growth in social media traffic
- Rollout of 5G services, enabled by massive small cell densification
- Accelerating deployments of IoT and IIoT (Industrial IoT)
- A shift from traditional office-based work to remote options

### Growth of hyperscale providers

Globally, true hyperscale data centers may number less than a dozen or so, but their impact on the overall data center landscaping is significant. According to recent research, the world spent a combined 1.25 billion years online in 2020 alone.<sup>1</sup> About 53% of that traffic passes through a hyperscale facility.<sup>2</sup>

### Hyperscale partnering with Multi-Tenant Data Center (MTDC / co-Location) facilities

As demand for more lower latency performance increases, hyperscale and cloud-scale providers work to extend their presence closer to the end user/end device. Many are partnering with MTDC or co-location data centers to locating their services at the so-called network "edge"<sup>3</sup>. When the edge is physically close, lower latency and network costs expand the value of new low-latency services. As a result, growth in the hyperscale arena is forcing MTDCs and co-location facilities to adapt their infrastructures and architectures to support the increased scale and traffic demands that are more typical of hyperscale data centers. At the same time, these largest of data centers must continue to be flexible to customer requests for cross-connections to cloud provider on-ramps.

## Spine-leaf and fabric mesh networks

The need to support low-latency, high-availability, very high bandwidth applications is hardly limited to hyperscale and co-location data centers. All data center facilities must now rethink their ability to handle the rising demands of end users and stakeholders. In response, data center managers are rapidly moving toward more fiber-dense mesh fabric networks. The any-to-any connectivity, higher fiber-count backbone cables and new connectivity options enable network operators to support ever-higher lane speeds as they prepare to make the transition to 400 Gigabits per second<sup>4</sup> (G).

## Enabling artificial intelligence (AI) and machine learning (ML)

In addition, the larger data center providers, driven in part by IoT and smart city applications, are turning to AI and ML to help create and refine the data models that power near real-time compute capabilities at the edge. Besides having the potential to enable a new world of applications (think commercially viable self-driving cars), these technologies require massive data sets, often referred to as data lakes, and massive compute power within the data centers and large enough pipes to push the refined models to the edge when needed.<sup>5</sup>

## Timing the move to 400G/800G


Just because you're running at 40G or even 100G today, don't be lulled into a false sense of security. If the history of data center evolution has taught us anything, it's that the rate of change—whether it's bandwidth, fiber density or lane speeds—accelerates exponentially. The transition to 400G is closer than you think. Not sure? Add up the number of 10G (or faster) ports you're currently supporting and imagine they progress to 100G, you'll realize that the need for 400G (and beyond) isn't that far away.

### What does it mean for data centers?


As data center managers look to the horizon, the signs of a cloud-based evolution are everywhere.



**More**  
high-performance  
virtualized servers




**Higher**  
bandwidth and lower  
latency



**Faster**  
switch-to-server  
connections



**Higher**  
uplink/backbone  
speeds



**Rapid**  
expansion capabilities

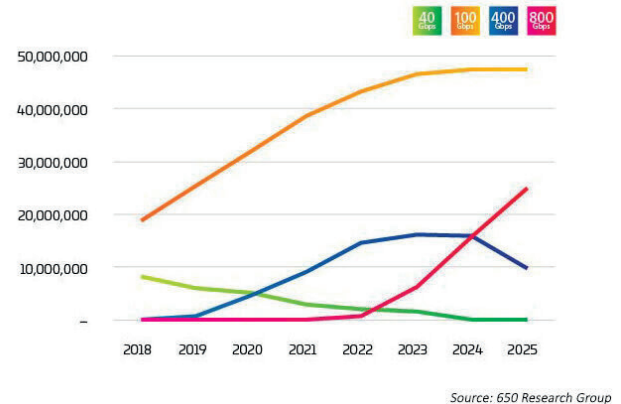
Within the cloud itself, the hardware is changing. Multiple disparate networks typical in a legacy data center have evolved to a more virtualized environment that uses pooled hardware resources and software-driven management. This virtualization is driving the need to route application access and activity in the fastest possible way, forcing many network managers to ask, "How do I design my infrastructure to support these cloud-first applications?"

The answer begins with enabling higher per-lane speeds. The progression from 25 to 50 to 100G and above is key to getting to 400G and beyond, and it has begun to replace the traditional 1/10G migration path. But there's more to it than increasing lane speeds, a lot more. We have to dig a bit deeper.

## Who (and what) is driving the change to higher speeds?

The industry is reaching an inflection point. The adoption of 400G has ramped up very quickly, but shortly 800G is expected to begin ramping even faster than 400G. As one might expect, there is no simple answer to “who or what is driving the transition to 400G?” There are a variety of factors at play, many of which are intertwined. New technologies enable lower cost per bit when lane rates increase. The latest data projects that 100G lane rates will be combined with octal switch ports to bring 800G options to market beginning in 2022. These ports are being utilized in several ways, however, as illustrated in the Light Counting data<sup>6</sup> where 400G and 800G are primarily broken out to 4X or 8X 100G. It is this breakout application that is the early driver of these new optic applications.

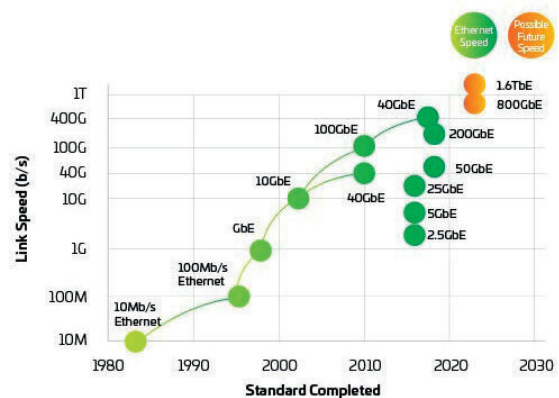
Figure 2: Data center Ethernet port shipments



## The transition to 400G is a group effort

In the data network, capacity is a matter of checks and balances among servers, switches and connectivity. Each pushes the other to be faster and less expensive, to efficiently track the demand produced by increased data sets, AI and ML. For years, switch technology was the primary bottleneck. With the introduction of Broadcom’s StrataXGS<sup>®</sup> Tomahawk<sup>®</sup> 3, data center managers can now boost switching and routing speeds to 12.8 Terabits/sec (Tb/s) and reduce their cost per port by 75 percent. Broadcom’s Tomahawk 4 switch chip, with a bandwidth of 25 Tb/s., provides the data center industry with more switching capability to stay ahead of those increasing AI and ML workloads. Today, this chip supports 64x 400G ports; but at 25.6Tb/s capacity, semiconductor technology is taking us down a path where in the future we could see 32x 800G ports on a single chip. 32, coincidentally, being the maximum number of QSFP-DD, or OSFP (800G transceivers) that can be presented at a 1U switch faceplate.

Figure 3: Ethernet speeds




So now, the limiting factor is the CPU processing capability. Right? Wrong. Earlier this year, NVIDIA introduced its new Ampere chip for servers. It turns out, the processors used in gaming are perfect for handling the training and inference-based processing needed for AI and ML. According to NVIDIA, one Ampere-based machine can do the work of 120 Intel-powered servers.

With switches and servers on schedule to support 400G and 800G by the time they’re needed, the pressure shifts to the physical layer to keep the network balanced. IEEE 802.3bs, approved in 2017, paved the way for 200G and 400G Ethernet. However, the IEEE has just now completed its [bandwidth assessment regarding 800G and beyond](#). The IEEE has started a study group to identify the objectives for applications beyond 400G and, given the time required to develop and adopt new standards, we may already be falling behind. The industry is now working together introducing 800G and beginning to work toward 1.6T and beyond while improving the power and cost per bit.

# The Four Pillars of 400G/800G migration

As you begin to consider the “nuts and bolts” of supporting your migration to 400G, it is easy to become overwhelmed by all the moving parts involved. To help you better understand the key variables that need to be considered, we’ve grouped them in four main areas:

- Increasing switch-port densities
- Optical transceiver technologies
- Connector options
- Cabling advances

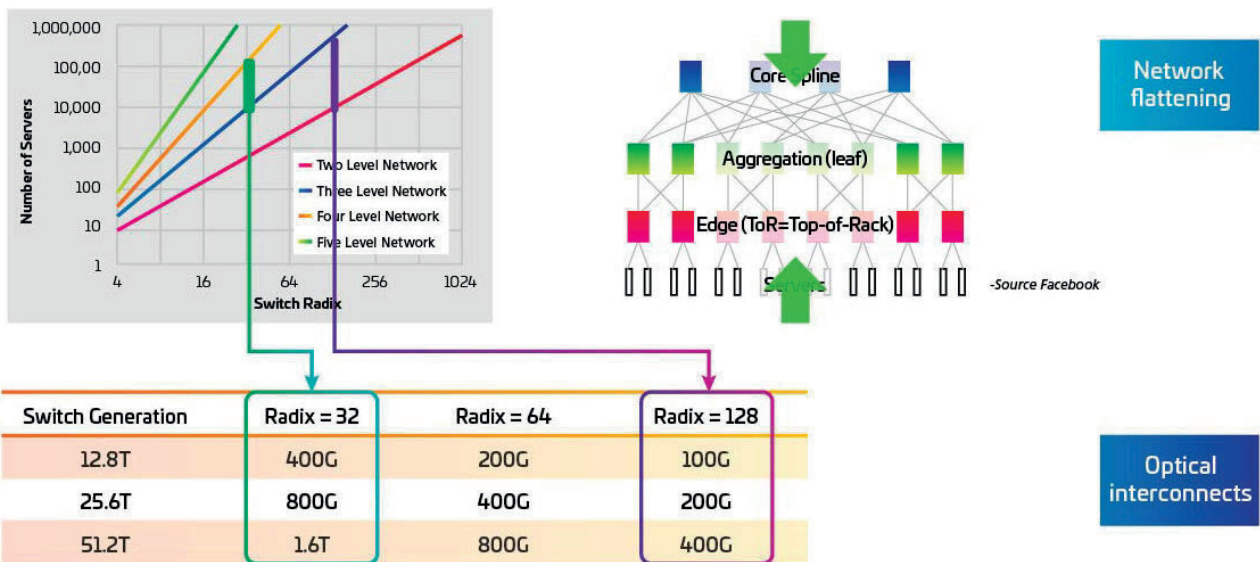


Together, these four areas represent a big part of your migration toolbox. Use them to fine-tune your migration strategy to match your current and future needs.

## Higher switch-port densities (radix) – The end of ToR?

Switching speeds are increasing as serializer/deserializer (SERDES) that provide the electrical I/O for the switching ASIC move from 10G, 25G, 50G. SERDES are expected to hit 100G once IEEE802.3ck becomes a ratified standard. Switch application-specific integrated circuits (ASICs) are also increasing the I/O port density (a.k.a. radix). Higher radix ASICs support more network device connections, offering the potential to eliminate a layer top-of-rack (ToR) switches. This, in turn, reduces the overall number of switches needed for a cloud network. (A data center with 100,000 servers can be supported with two levels of switching with a RADIX of 512.) The higher radix ASICs translate into lower CAPEX (less switches), lower OPEX (less energy required to power and cool fewer switches) and improved network performance through lower latencies.

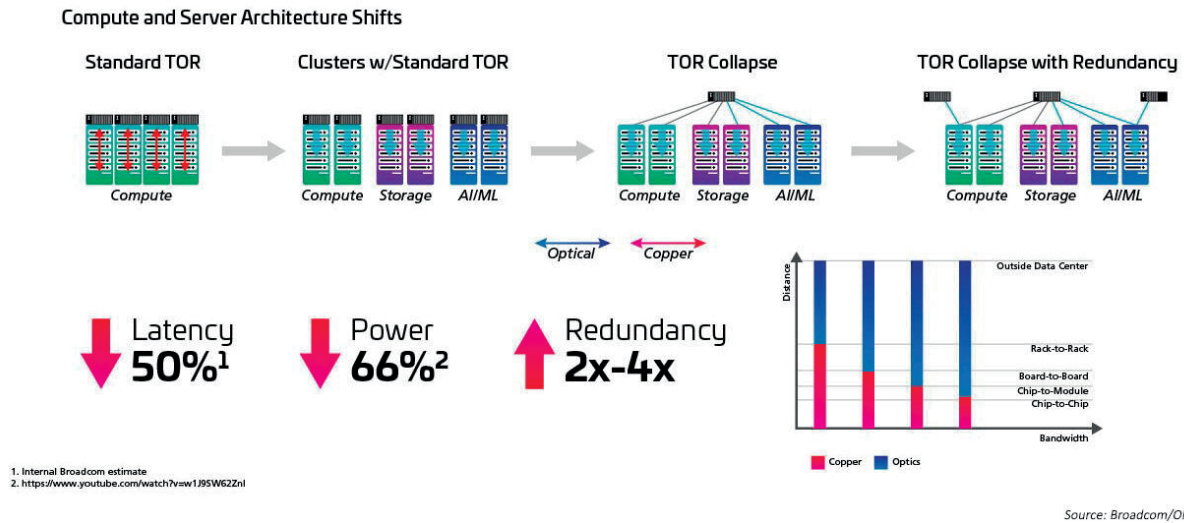
**Figure 4: Effects of higher Radix switches on switch bandwidth**



Source: OFC 2021, Mark Filer Microsoft on Lyubomirsky ECOC'20 importance of Radix

Closely related to the increase in radix and switching speed is the move from a top-of-rack (ToR) topology to either middle-of-row (MoR) or end-of-row (EoR) configuration, and the benefit that structured cabling approach holds when facilitating the many connections between the in-row servers and the MoR/EoR switches. The ability to manage the large number of server attachments with greater efficiency is required to make use of new high-radix switches. This, in turn, requires new optic modules and structured cabling, such as those defined in the IEEE802.3cm standard. The IEEE802.3cm standard supports the benefits of pluggable transceivers for use with high-speed server network applications in large data centers defining eight host attachments to one QSFP-DD transceiver

**Figure 5: Architectures shifting from ToR to MoR/EoR**



## Transceiver technologies

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. No problem.

QSFP-Double Density (QSFP-DD<sup>7</sup>) and octal (2 times a quad) small form factor pluggable (OSFP<sup>8</sup>) 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 when attaching servers.


## New transceiver formats

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.

The “second-generation” 400G form factor modules feature QSFP-DD and OSFP. The QSFP-DD transceivers are backwards compatible with existing QSFP ports. They build on the success of the existing optic modules, QSFP+ (40G), QSFP28 (100G) and QSFP56 (200G).

OSFP, like the QSFP-DD optics, enables use of eight lanes versus four. Both types of modules support 32 ports in a 1RU card (switch). To support backwards compatibility, the OSFP requires an OSFP-to-QSFP adapter.

Figure 6: OSFP versus QSFP-DD transceiver



### QSFP-DD (type 1)

Dimensions (mm)  
18.35 x 58.2 x 8.5  
Volume (mm<sup>3</sup>)  
9,087

### OSFP

Dimensions (mm):  
22.93 X 100.4 X 13  
Volume (mm<sup>3</sup>)  
29,928

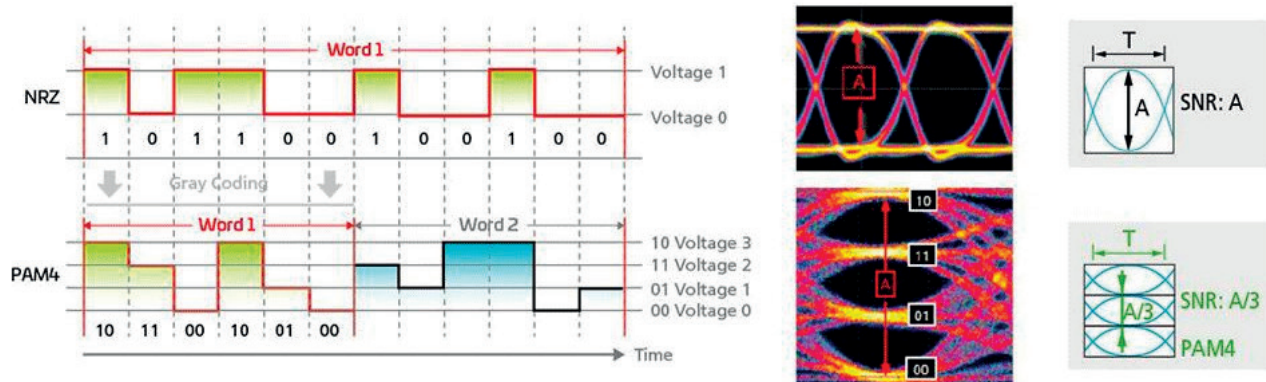
Source: Fluxlight

## Modulation schemes

Network engineers have long utilized non-return to zero (NRZ) modulation for 1G, 10G and 25G, using host-side forward error correction (FEC) to enable longer distance transmissions. To get from 40G to 100G, the industry simply turned to parallelization of the 10G/25G NRZ modulations, also utilizing host side FEC for the longer distances. When it comes to achieving speeds of 200G/400G and faster, new solutions are needed.

Figure 7: Higher-speed modulation schemes are used to enable 50G and 100G technologies

### NRZ and PAM4 Encoding



Source: Samtec

As a result, optical networking engineers have resorted to four-level pulse amplitude modulation (PAM4) to bring ultra-high bandwidth network architectures to fruition; PAM4 is the current solution for 400G. This is based in large measure on IEEE802.3, which has completed new Ethernet standards for rates up to 400G (802.3bs/cd/cu) for both multi-mode (MM) and single-mode (SM) applications. A variety of breakout options are available to accommodate diverse network topologies in large scale data centers.

More complex modulation schemes imply the need for an infrastructure that can provide better return loss and attenuation.



## 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 favored for DCI optical links or those specifically requiring higher power and more optical I/Os. OSFP proponents envision 1.6T and perhaps 3.2T transceivers in the future.

Co-packaged optics (CPOs) provide an alternate path to 1.6T and 3.2T. But CPOs will need a new ecosystem that can move the optics closer to the switch ASICs to achieve the increased speeds while reducing power consumption. This track is being developed in the Optical Internetworking Forum (OIF). The OIF is now discussing the technologies that might be best suited to the "next rate," with many arguing for a doubling to 200G. Other options include more lanes – perhaps 32, as some believe that more lanes and higher lane rates will eventually be needed to keep pace with network demand at an affordable network cost.

The only sure prediction is that the cabling infrastructure must have the built-in flexibility to support your future network topologies and link requirements. While astronomers have long held that "every photon counts" as network designers look to reduce the energy per bit to a few pJ/Bit<sup>9</sup>, conservation at every level is important. High-performance cabling will help reduce network overhead.



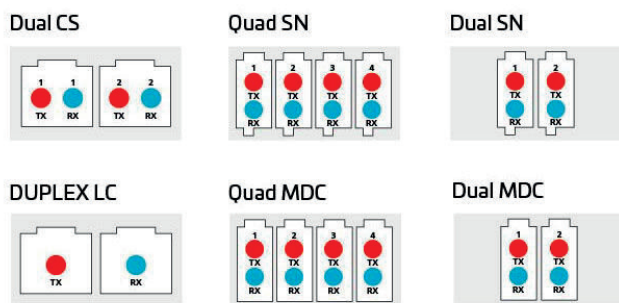
Figure 8: Options for distributing capacity from octal modules

### MPO and duplex connectors

Switches are evolving to provide more lanes at higher speeds while reducing the cost and power of networks. Octal modules allow these additional links to connect through the 32-port space of a 1U switch. Maintaining the higher radix is accomplished by using lane breakout from the optic module.

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

Figure 8: Options for distributing capacity from octal modules



Source: QSFP-DD/QSFP-DD800/QSFP112 Hardware Specification

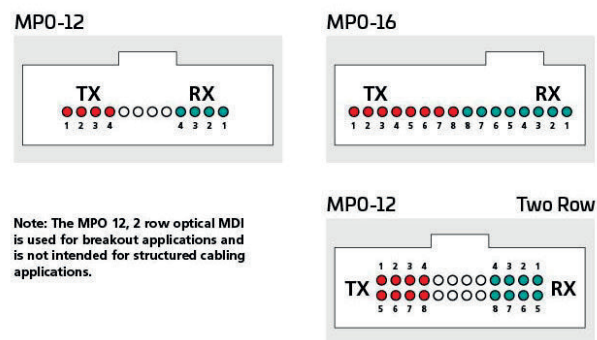
## MPO connectors

Until recently, the primary method of connecting switches and servers within the data center involved cabling organized around 12- or 24-fibers, typically using MPO connectors. The introduction of octal technology (eight switch lanes per switch port) enables data centers to match the increased number of ASIC I/Os (currently 256 per switch ASIC) with optical ports. This yields the maximum number of I/Os available to connect servers or other devices.

The optical I/Os use connectors that are appropriate for the number of optical lanes used. A 400G transceiver may have a single duplex LC connector with a 400G optical I/O, it could also have 4 X 100G optical I/Os requiring 8 fibers. The MPO12 or perhaps 4 SN duplex connectors will fit within the transceiver case and provide the 8 fibers this application needs. Sixteen fibers are required to match 8 electrical and optical I/Os, preserving the radix of the switch ASIC. The optic ports can be either single mode or multimode depending on the distance the link is designed to support.

For example, multimode technology continues to provide the most cost-effective high-speed optical data rates for short reach links in the data center. IEEE standards support 400G in a single link (802.3 400G SR4.2) technology, which uses four fibers to transmit and four fibers to receive, with each fiber carrying two wavelengths. This standard extends the use of bi-directional wavelength division multiplexing (BiDi WDM) techniques and was originally intended to support switch-to-switch links. This standard uses the MPO12 connector and was the first to optimize using **OM5 MMF**.

Figure 9: MPO 12- and 16-fiber configurations



Source: QSFP-DD/QSFP-DD800/QSFP112 Hardware Specification

Maintaining the switch radix is important where many devices, such as server racks, need to be connected to the network. 400G SR8, addressed in the IEEE 802.3cm standard (2020), supports eight server connections using eight fibers to transmit and eight fibers to receive. This application has gained support amongst cloud operators. MPO-16 architectures are being deployed to optimize this solution. Single-mode standards support longer reach applications (switch-to-switch, for example). IEEE 400G-DR4 supports 500 meters reach with 8 fibers. This application can be supported by MPO-12 or MPO-16. The value of the 16-fiber approach is added flexibility; data center managers can divide a 400G circuit into manageable 50/100G links. For example, a 16-fiber connection at the switch can be broken out to support up to eight servers connecting at 50/100G while matching the electrical lane rate. MPO 16-fiber connectors are keyed differently to prevent connection with the 12-fiber MPO connectors.

The electrical lane rate then determines the output capabilities of the optical interface. Table 1 shows examples of the 400G (50G X 8) module standards/possibilities.

**Table 1: 400G Capacity QSFP-DD with 50G Electrical Lanes**

Application				Module Naming Schemes		
Reach	4x100G Ethernet	2x200G Ethernet	400G Ethernet	Scheme A	Scheme B	Scheme C
SR (50-70m)	4x100G-BiDi 4x100GBASE-SR2 4x100GBASE-SR1	2x200GBASE-SR4 2x200GBASE-SR2	400GBASE-SR4.2 400GBASE-SR8 400GBASE-SR4	QDD-400G-SR4.2 QDD-400G-SR8* QDD-400G-SR4	400G-BiDi 400G-SR8 400G-SR4	
DR (500m)	4x100GBASE-DR	2x200GBASE-DR4	400GBASE-DR4	QDD-2x200G-DR4 QDD-400G-DR4	400G-DR8 400G-DR4	
FR (2km)	4x100GBASE-FR1	2x200GBASE-FR4	400GBASE-FR8 400GBASE-FR4	QDD-4x100G-FR1 QDD-2x200G-FR4 QDD-400G-FR8 QDD-400G-FR4	400G-4xFR1 400G-2xFR4 400G-FR8 400G-FR4	400G-DR4+
LR (6km)	-no standard-	-no standard-	400GBASE-LR4-6	QDD-400G-LR4-6	400G-LR4-6	
LR (10km)	4x100GBASE-LR1	2x200GBASE-LR4	400GBASE-LR8 400G-LR4-10	QDD-4x100G-LR1 QDD-2x200G-LR4 QDD-400G-LR8 QDD-400G-LR4-10	400G-4xLR1 400G-2xLR4 400G-LR8 400G-LR4-10	400G-DR4++
ER (30-40km)			400GBASE-ER8 400ZR (unamp)	QDD-400G-ER8 QDD-400ZR	400G-ER8 400ZR	
ZR (80-120km)			400ZR(amp) 400GBASE-ZR	QDD-400ZR QDD-400G-ZR	400ZR 400G-ZR	

\*QDD-400G-SR8 also supports 8 x 50GBASE-SR

When lane rates are doubled to 100G, the following optical interfaces become possible. At the time of writing, 100G lane rate standards (802.3 ck) have not been completed; however, early products are being released and many of these possibilities are in fact shipping. Table 2, presented at ECOC 2020 by J. Maki (Juniper), shows the early industry interest in the 800G modules.


**Table 2: 800G Capacity QSFP-DD with 100G Electrical Lanes**

Application					Module Naming Schemes	
Reach	8x100G Ethernet	4x200G Ethernet	2x400G Ethernet	800G ETC Eth.	Scheme A	Scheme B/C
SR (50m)	8x100GBASE-SR1	4x200GBASE-SR2	2x400GBASE-SR4	No ask	QDD-2x400G-SR4	800G-SR8
DR (500m)	8x100GBASE-DR		2x400GBASE-DR4	No ask	QDD-2x400G-DR4	800G-DR8
FR (2km)	8x100GBASE-FR1		2x400GBASE-FR4	800G-FR4/FR8	QDD-8x100G-FR1 QDD-2x400G-FR4 QDD-800G-FR4/FR8	800G-DR8+ 800G-2xFR4 800G-FR4/FR8
LR (6km)			2x400GBASE-LR4-6	No ask	QDD-2x400G-LR4-6	800G-2xLR4-6
LR (10km)	8x100GBASE-LR1		2x400G-LR4-10	No ask	QDD-8x100G-LR1 QDD-2x400G-LR4-10	800G-DRB++ 800G-2xLR4-10
LR (high loss)				OIF 800G Coherent	QDD-800LR	800LR
ER (30-40km)	No ask		No ask	No ask	-	-
ZR (80-120km)	No ask		No ask	OIF 800G Coherent	QDD-800ZR	800ZR

## Duplex connectors

As the number of lanes and lane speeds increase, splitting the optical I/Os becomes more attractive. As mentioned earlier, octal modules can support connector options for 1, 2, 4 or 8 duplex links. All these options can be accommodated using an MPO connector; however, that option may not be as efficient as separate duplex connectors. A duplex connector with a smaller footprint can help enable these options. The SN, a very small form factor (VSFF) duplex optical fiber connector, fits this application. It incorporates the same 1.25mm ferrule technology previously used in the LC connectors. As a result, it delivers the same optical performance and strength but is targeted to more flexible breakout options for high-speed optical modules. The SN connector can provide four duplex connections to an octal transceiver module. The early applications for the SN are primarily to enable optical module breakout applications.

**Figure 10: Size relationship between leading duplex connectors and breakout applications for 400G/800G migration**



## Connector speed limits?

Connectors typically do not dictate speed, economics do. Optical technologies were initially developed and deployed by service providers who had the financial means and bandwidth demands to support their development, as well as the long-haul links that are most economically bridged using the fewest number of fibers. Today, most service providers prefer simplex or duplex connector technology paired with optical transport protocols that use single-fiber connector technologies like LC or SC.

However, these long-haul solutions can be too expensive, especially when there are hundreds or thousands of links and shorter link distances to traverse; both conditions are typical of a data center. Therefore, data centers often deploy parallel optics. Since parallel transceivers provide a lower cost per Gigabit, MPO-based connectivity is a good option over shorter distances. Thus, connector choices today are not driven so much by speed, but by the number of data lanes they can support, the space they take up, and the price impact on transceivers and switch technologies.

In the final analysis, the range of optical transceivers and optical connectors is expanding, driven by a wide variety of network designs. Hyperscale data centers may choose to implement a very custom optical design; given the scale of these market movers, standards bodies and OEMs often respond by developing new standards and market opportunities. As a result, investment and scale lead the industry in new directions and cabling designs evolve to support these new requirements.

### Cabling advances

This topic will be the subject for Part II. Stay tuned

# Further information

## Additional helpful DC resources

### Additional helpful DC resources



**Solution**  
Enterprise data centers

[Learn more](#)



**Solution**  
Hyperscale and Cloud  
Data Centers

[Learn more](#)



**Solution**  
Multi tenant data  
centers

[Learn more](#)




**Solution**  
Service provider data  
centers

[Learn more](#)



**Insights**  
Multimode Fiber: the  
Fact File

[Read](#)



**Resources**  
High Speed Migration  
Library

[Learn more](#)



Source: osfpmsa.org

**Specification info**  
OSFP MSA


[Learn more](#)



Source: qsfp-dd.com

**Specification info**  
QSFP-DD MSA

[Learn more](#)



Source: qsfp-dd.com

**Specification**  
QSFP-DD Hardware

[Learn more](#)

## Why choose CommScope?

At first glance, the field of potential infrastructure partners vying for your business seems pretty crowded. There's no shortage of providers willing to sell you fiber and connectivity. But as you look closer and consider what's critical to the long-term success of your network, the choices begin to narrow. That's because it takes more than fiber and connectivity to fuel your network's evolution...a lot more. That's where CommScope stands out.



**Proven Performance:** CommScope's history of innovation and performance spans 40+ years—our singlemode TeraSPEED® fiber debuted three years before the first OS2 standard, and our pioneering wideband multimode gave rise to OM5 multimode. Today, our end-to-end fiber and copper solutions and AIM intelligence support your most demanding applications with the bandwidth, configuration options and ultra-low loss performance you need to grow with confidence.



**Agility and adaptability:** Our modular portfolio enables you to quickly and easily respond to shifting demands in your network. Singlemode and multimode, pre-terminated cable assemblies, highly flexible patch panels, modular components, 8-, 12-, 16- and 24-fiber MPO connectivity, very small form factor duplex and parallel connectors. CommScope keeps you fast, agile and opportunistic.



**Future-ready:** As you migrate from 100G to 400G, 800G and beyond, our high-speed migration platform provides a clear, graceful path to higher fiber densities, faster lane speeds and new topologies. Collapse network tiers without replacing the cabling infrastructure, move to higher-speed, lower-latency server networks as your needs evolve. One robust and agile platform takes you from now to next.



**Guaranteed reliability:** With our Application Assurance, CommScope guarantees that the links you design today will meet your application requirements years down the road. We back that commitment with a holistic lifecycle service program (planning, design, implementation and operation), a global team of field application engineers and CommScope's iron-clad 25-year warranty.



**Global availability and local support:** CommScope's global footprint includes manufacturing, distribution and local technical services that span six continents and features 20,000 passionate professionals. We are there for you, whenever and wherever you need us. Our global Partner Network ensures you have the certified designers, installers and integrators to keep your network moving forward.



High Speed Migration



SYSTIMAX® Fiber Structured Cabling



Copper Trunk Assemblies



Fiber Raceways



Automated Management

1 Digital Trends 2020; thenextweb.com

2 The Golden Age of HyperScale; Data Centre magazine; November 30, 2020

3 [https://attom.tech/wp-content/uploads/2019/07/TIA\\_Position\\_Paper\\_Edge\\_Data\\_Centers.pdf](https://attom.tech/wp-content/uploads/2019/07/TIA_Position_Paper_Edge_Data_Centers.pdf)

4 <https://www.broadcom.com/blog/switch-phy-and-electro-optics-solutions-accelerate-100g-200g-400g-800g-deployments>

5 The Datacenter as a Computer Designing Warehouse-Scale Machines Third Edition Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan Google LLC. Morgan & Claypool publishers pg 27

6 [LightCounting presentation for ARPA-E conference - October 2019.pdf](#) (energy.gov)

7 <http://www.qsfp-dd.com/wp-content/uploads/2021/05/QSFP-DD-Hardware-Rev6.0.pdf>

8 [https://osfpm.org/assets/pdf/OSFP\\_Module\\_Specification\\_Rev3\\_0.pdf](https://osfpm.org/assets/pdf/OSFP_Module_Specification_Rev3_0.pdf)

9 Andy Bechtolsheim, Arista, OFC '21

## The Move to 400Gb is closer than you think

Get the shortlist of the signs that point out to the cloud-based DC evolution.

[Read](#)



CommScope pushes the boundaries of communications technology with game-changing ideas and ground-breaking discoveries that spark profound human achievement. We collaborate with our customers and partners to design, create and build the world's most advanced networks. It is our passion and commitment to identify the next opportunity and realize a better tomorrow. Discover more at [commscope.com](https://commscope.com)

**COMMSCOPE®**

---

[commscope.com](https://commscope.com)

Visit our website or contact your local CommScope representative for more information.

© 2021 CommScope, Inc. All rights reserved.

Unless otherwise noted, all trademarks identified by ® or ™ are registered trademarks, respectively, of CommScope, Inc. This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services. CommScope is committed to the highest standards of business integrity and environmental sustainability with a number of CommScope's facilities across the globe certified in accordance with international standards, including ISO 9001, TL 9000, and ISO 14001. Further information regarding CommScope's commitment can be found at [www.commscope.com/About-Us/Corporate-Responsibility-and-Sustainability](https://www.commscope.com/About-Us/Corporate-Responsibility-and-Sustainability)

MC-115887-EN (12/21)