High Speed Migration:

Choosing the right multimode multi-fiber push on (MPO) system for your data center

Table of Contents

In the quest to deliver ever higher speeds, the infrastructure for data center networks, large and small, is increasingly dominated by optical fiber. Traditionally, multimode and singlemode optical networks have been based on duplex fiber links. Parallel fiber (MPO-based) connectivity has become popular, since it enables the use of preterminated systems that can be deployed quickly and efficiently, but it has been predominantly used to deliver duplex connectivity in combination with duplex modules and breakouts. More recently, however, standards and equipment vendors have begun to deploy multi-fiber interfaces to increase throughput and density.

These developments are leading customers to consider not only the delivery of duplex connectivity, but to plan for a high-speed migration strategy that may involve a mix of parallel and duplex interfaces, a mix that may change as new applications are deployed. To define the optimal strategy, it is important to understand the expected evolution of network equipment, as well as the benefits and tradeoffs of deploying an infrastructure based on duplex connections and/or 8-fiber, 12-fiber and 24-fiber MPOs.

Many types of applications coexist in enterprise data centers. Data networks interconnect storage devices to application servers —sometimes storage traffic is converged with other Ethernet IP network resources. Applications often require access to different resources in the data center and this drives a variety of connectivity topologies. Today, most network links in enterprise data centers utilize a single pair of fibers (duplex links). However higher speed with more capacity can be achieved by combining duplex links into groups of links (parallel links). Cabling designs must support both design options and the migration that occurs from duplex to parallel and back again over time. There are many examples of this trend. Fibre Channel applications for storage area networks have remained primarily duplex, but as we migrate beyond 32G FC new parallel link options are being introduced. We have witnessed a similar transition in Ethernet networks. Some Ethernet applications are evolving from duplex to parallel and then back again to duplex. For example, multimode 40G and 100G Ethernet were initially enabled with 4 x 10G and 10 x 10G duplex channels respectively. Today, 100G is primarily delivered on 4 x 25G duplex pairs instead of ten, and both 40G and 100G can now be deployed over a single duplex fiber pair. However, there are other reasons that parallel optics are used. In network equipment, the popular 8-fiber parallel quad small form-factor pluggable (QSFP) port provides four times the connectivity density and reduces the power consumption compared to equivalent SFP interfaces by approximately 30 percent. This lowers the initial capital cost of the network as well as the ongoing operating costs. At the other end, these parallel ports are typically “broken out” as duplex connections, providing connectivity to four end devices.

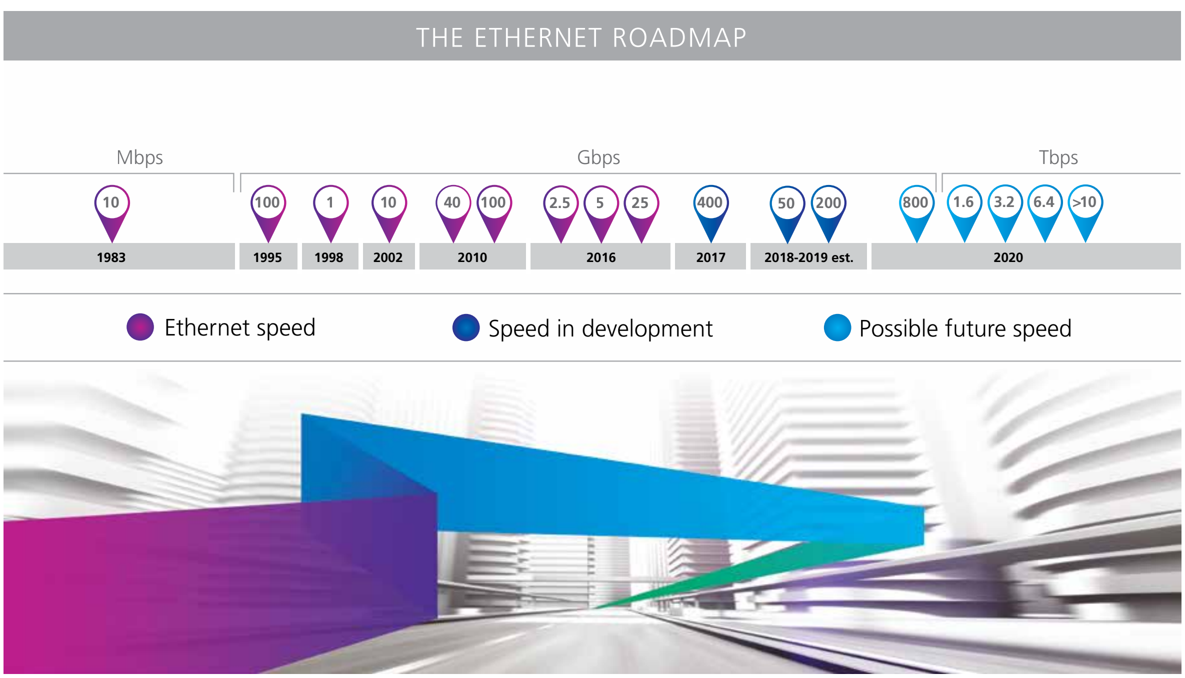

Figure 1: The Ethernet road map

The Ethernet march to higher speeds is forcing interface designers to break from the tradition of 10x increases in speed for every iteration to a modular and more gradual approach that focuses on the individual transmission "lanes." These lanes may be combined to increase throughput by using additional fibers, multiplexing over a single fiber, or both. Starting with 40G, which was initially achieved by combining 4 x 10G lanes over eight fibers, Institute of Electrical and Electronics Engineers (IEEE) specifications evolved to include the initial 100G specification delivering 10 x 10G lanes over 20 fibers, before moving up to 25G and 50G lanes as building blocks. When first moving to higher speeds, parallel applications are often the most economical and shortest path. As the technology progresses, implementing higher speeds on a single pair of fibers becomes cost effective, as we have seen in the past with 40G Ethernet and most recently, with 100G Ethernet (4 x 25G lanes). As 50G lanes become more cost effective and popular the process will continue—both enabling higher parallel link speeds like 200G and 400G, but at the same time, enabling higher duplex fiber speeds by deploying other techniques such as wave division multiplexing.

Designing fiber optic cabling for data centers requires a strategy to balance the cost/benefit of the opto-electronics against the investment choices for fiber optic cabling. One basic trade-off is a choice between parallel and duplex fiber links. This choice can optimize the initial investment in cabling infrastructure and optical transceivers against the longer term future requirements. If the data center demand is growing quickly and the life cycle of compute/storage and network equipment is relatively short, your strategy to support these changes might suggest a preference for duplex or parallel optics, and that will influence the cabling choices. If you are unsure of the future path, then looking for options that can provide optimal support for any of these potential outcomes would be favored.

This paper will provide a framework to evaluate and choose the fiber media and cabling options that best suit your data center delivery strategy. These key design elements of fiber optic infrastructures should be standard:

- Enabling rapid deployment of new resources and higher speed networks
- Full support for parallel or duplex optics and a clear path forward for high speed migration throughout the planning period.
- High performance optical components including preterminated MPO-based systems

Planning considerations

As a data center infrastructure designer, the capacity that your data center will require in the future and the time lines for managing these changes are key data points to consider. Certainly, there are many other considerations that can help guide design choices such as:

- Is your organization an early adopter of new technology?
- Is the growth of the organization slow/fast or unpredictable?
- What are the capital vs operational cost tradeoffs?
- Service for internal networks, multi-tenant data center (MTDC), private/public cloud?
- Are your applications and the compute/storage structures changing?

The answers to questions such as these will help guide the design choices. Quite often the fiber infrastructure designer simply does not have all the data points they would like to complete their plans. Data center teams must evaluate many different technologies to complete a new design strategy. Server teams must evaluate various platforms, such as server attached storage vs. Fibre Channel, etc. Network teams must evaluate various fabric and routing/switching strategies. All of these technologies are quickly changing. Even in the absence of all the facts we would like to have on hand, the fiber infrastructure can be designed to meet the outcome of the data center plans. When properly selected and provisioned, preterminated MPO-based fiber cabling systems enable rapid deployment and configuration flexibility. They form the basis of a strategy to quickly deliver cabling topologies that support the data center direction, whatever that turns out to be, with optimized deployment costs.

Fiber media

Multimode fiber is the primary media choice for enterprise data centers. There are various types of multimode fiber (MMF) available and the characteristics of the MMF will determine the scale and scope of data center that can be supported when speeds escalate. It will also have a bearing on when duplex optical technology will be available compared to parallel options for a given speed. Some large enterprise data centers are now deploying single mode optics and infrastructure in their data centers. As is typical with early stages of optic technologies these are, for the most part, being implemented with parallel optics/fiber. Not all cabling options available for MMF are available for singlemode fiber (SMF) but, generally, the principles of the cabling designs we will discuss will apply to both MMF and SMF fiber.

Network optics

We discussed the provisioning of multiple "lanes" that, when taken together, provide for greater link speeds, as was the case with 40GBASE-SR4. Another method to increase the number of lanes is to multiplex several wavelengths on a single pair of fibers. Two examples of this short wave division multiplexing (SWDM) technology are shown in figure 2, as 40G-SWDM4 and 100G-SWDM4. These transceivers provide four 25G lanes over a duplex pair of multimode fibers using four different wavelengths. In support of these new technologies, a new OM5 multimode media was introduced to provide extended support for these additional wavelengths. Going forward, 100G-SWDM duplex can then be aggregated into a parallel 400G link using 8 fibers. Note the fiber counts for common applications in the chart below. Any of the multi-pair applications shown can break-out too, so we expect to see migrations between duplex and parallel optics and, therefore, expect that fiber cabling systems should be provisioned to accommodate these changes in the future.

	Standard/ (# fibers)	Maximum distance
40G	**40GBASE-SR4 (8)**	OM3 100m OM4/OM5 150m
40G	400G-BiDi (2)	OM3 100m* OM4/OM5 150m* OM5 200m
40G	40GBASE-eSR4 (8)	OM3 300m OM4/OM5 400m
40G	40G-SZDM4 (2)	OM3 240m* OM4 350m* OM5 440m
40G	**100GBASE-SR4 (8)**	OM3 70m OM4/OM5 100m
100G	**100GBASE-SR10 (20)**	OM3 100m OM4/OM5 150m
100G	100GBASE-eSR4 (8)	OM3 200m OM4/OM5 300m
100G	100G-SWDM4 (2)	OM3 75m* OM4 100m* OM5 150m

Figure 2: Ethernet network optic application - * OM3/OM4 effective modal bandwidth is only specified at 850nm

Cabling density

The amount of space within a rack or cabinet dedicated to passive cabling is usually limited and very difficult to create or add after the initial installation. A two-fiber duplex LC port takes the same space as a single MPO port, but the MPO port can contain up to 24 fibers, providing much higher connection density. On the other hand, if you design only for parallel applications a change to duplex will require at least 4x the panel space in the rack or cabinet to accommodate the duplex LC ports. MPO-LC break-out assemblies can also convert MPO trunks to LC duplex connections if proper cable management techniques are used.

Duplex LC ports can be exchanged for MPO ports as needed. If the optics can be supported with a single duplex pair, then MPO ports can support up to twelve times as many links as LC ports. MPO connectors are available with 8, 12 or 24 fibers as illustrated in figure 3.

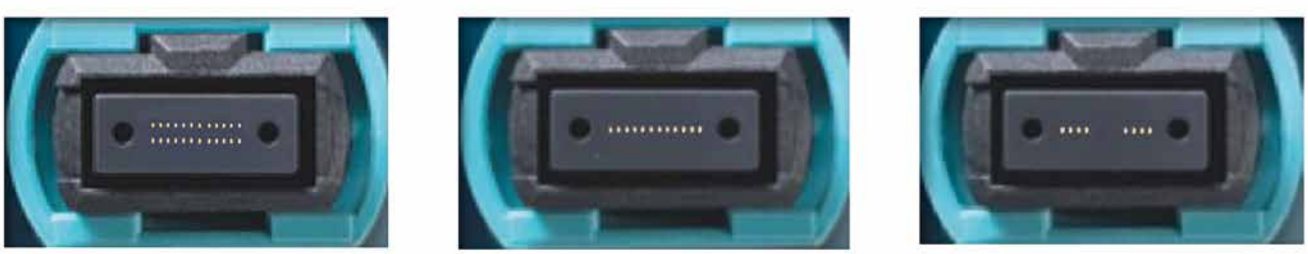

Figure 3: MPO-24, MPO-12 and MPO-8 configurations

Designing the fiber infrastructure with MPO systems

8-FIBER MPO (MPO-8) DESIGNS

The MPO-8 supports up to four duplex channels, as previously described, and is used in the QSFP transceivers that utilize the 8 fibers to provide four lanes per MPO port.

MPO-8 – connector

The MPO-8 interface is not strictly a standards recognized infrastructure interface. It is an application that happens to use 8 of the 12 fiber positions of the industry standard MPO-12. The QSFP application uses the outer 4 fiber locations (1-4 and 8-12) of the MPO-12 connector with 4 fibers used to transmit and 4 fibers used to receive. A QSFP transceiver does not use the center 4 fiber positions in the MPO-12 connector. While a standard MPO-12 cable will fully support this application, leaving fibers unused is not desirable. For this reason, cables and connectors can be constructed with MPO-12 connectors simply not providing or terminating the center four fibers as shown below in figure 4.

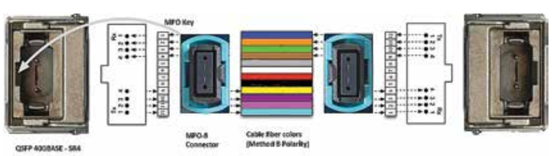

Figure 4: MPO-8 – QSFP applications with method B polarity management

In some cases, the QSFP application is used as a grouped link or trunk, between network switches. This application is supported with MPO-8, MPO-12 or MPO-24 systems. The number of fibers in the MPO connector determines how many ports the MPO trunk can support. In one case, MPO-8 connectors and trunks are used 1:1 for this application. MPO-12 trunks support this application in a 2:3 ratio while MPO-24 trunks support this application at 1:3 ratio.

In all MPO fiber systems, regardless of the number of fibers in the connector, polarity or routing a transmitted signal to the correct receiver is critical. While there are several polarity methods that may be used, method B as illustrated in figure 4 above and figure 5 below is the recommended option. This method uses aligned key adapters. MPO connectors use pins to ensure fibers are properly aligned, eliminating excess loss. Transceivers contain pins and the mating MPO cord connector are pinless. This pinning method is repeated throughout the link.

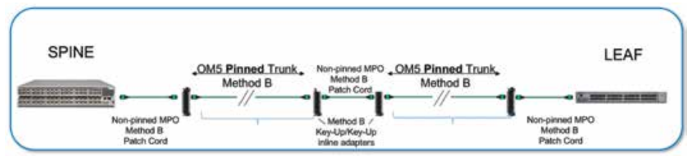

Figure 5: MPO-8 – Parallel cabling

The other common form of the QSFP application is in a breakout of one transceiver to four other devices as pictured in figure 6. Often this occurs within a cabinet or rack as is the case with top of rack to server connections. Occasionally, the switch and servers are located in different areas of the data center. In this case, the QSFP port is carried across the trunk and broken out at the server end of the trunk. It could also be broken out at the switch end and carried across the trunk as duplex ports and cabled one to one at the server end. As this is essentially a duplex application there is no compelling reason to use MPO-8 trunks to support it, and MPO-12 or MPO-24 trunks provide much more cost-effective support for duplex applications.

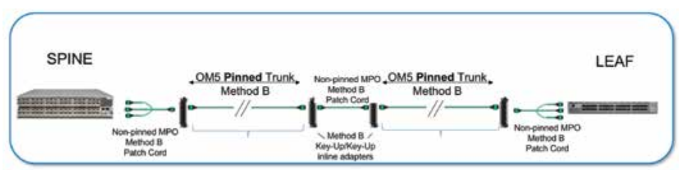

Figure 6: MPO-8 – QSFP breakout applications — 4 duplex ports

While it is possible to use MPO-8 trunks to support duplex applications as illustrated in figure 7, it requires at least 50 percent more connectors and panel space compared to MPO-12 or 3x the connectors and panel space compared to MPO-24. Additional connectors increase the capital cost and panel costs of this design. More connectors also increase the expense associated with cleaning and inspection of these connectors.

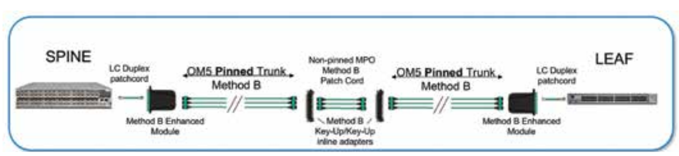

Figure 7: MPO-8 – Parallel cabling

12-FIBER (MPO-12) DESIGNS

The MPO-12 is a globally recognized standard interface for multimode and singlemode applications. The MPO-12 has been available for several decades and widely used as trunk cable connector for duplex and simplex applications.

MPO-12 – connector

The MPO-12 is often used as a duplex trunk as illustrated in figure 9. In the figure, two trunks are combined to support 24 fiber LC breakout modules (12 duplex ports). Again, the polarity control scheme that is preferred for this application is method B. The breakout modules can be varied with modules designed to provide breakout to MPO-8 (with 2 x MPO-12 ports and 3 x MPO-8 ports).

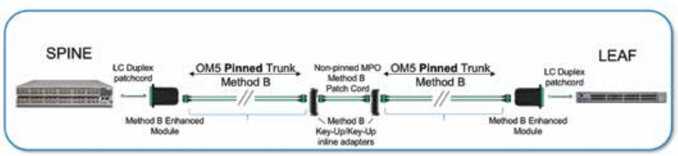

Figure 9: MPO-12 group of 2 trunks for 12 LC ports

Fan-out assemblies can take the place of modules and connect directly with MPO-12 trunks to provide the same functionality as the breakout modules as illustrated in figure 10. In some cases, breakout assemblies are preferred as they offer much higher panel density and reduce the number of connections in the network channel. This is balanced against the cable management requirements. Breakout assemblies and modules can be mixed in various combinations to suit the applications being supported.

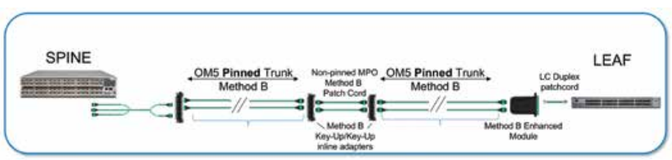

Figure 10: MPO-12 modules and fan-out assembly example

This example also introduces the ability of MPO-12 systems supporting 8-fiber parallel applications. This application presents a challenge, as the QSFP application does not make use of all twelve fibers as we have previously discussed. To maintain the efficient use of the fiber trunk cables the center four fibers are combined from the two 12-fiber trunks to provide an additional QSFP 8-fiber port.

Figure 11 illustrates the QSFP MPO-8 ports supported with MPO-12 trunks. In this case, 12-port duplex LC modules have been replaced with modules providing three QSFP ports. All of the 24 fibers in the MPO-12 trunks are fully utilized. Panel density is improved for cross connecting the trunks as only two MPO ports are used for three QSFP ports. Panel density for equipment interfaces is reduced as twelve LC port positions are replaced by three QSFP ports. This design does allow for switching back and forth between parallel and duplex applications without impacting the equipment panel space and with minimal differences to cord management requirements.

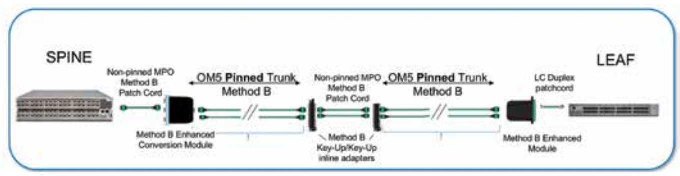

Figure 11: Three port QSFP module supported with two MPO-12 trunks

24-FIBER MPO (MPO-24) DESIGNS

The MPO-24 connector is perhaps the most cost effective method to deploy both parallel and duplex fiber optic applications. With twenty-four fibers in a single connector it provides additional density versus three MPO-8 connectors or two MPO-12 connectors, and expedites the cleaning and inspection time associated with the installation of the MPO systems. MPO-24 trunk cables are illustrated in figure 12. The method B trunk cables manage the port polarity in a similar fashion to the MPO-8 and MPO-12. The cross connection of MPO-24 trunks provides much higher port densities, reducing panel space requirements by 3:1 compared to MPO-8 and 2:1 compared to MPO-12. In high density applications, the trunk cable size is often a consideration as well. A 144-fiber trunk cable using MPO-24 subunits occupies approximately 30 percent less area than the MPO-12 equivalent.

MPO-24 – connector

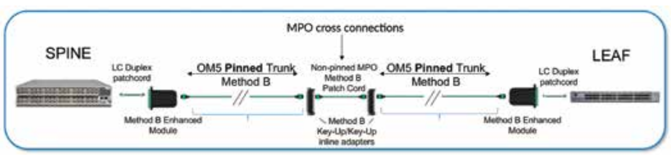

Figure 12: MPO-24 duplex applications

Figure 13 illustrates the MPO-LC break-out modules configured to provide cross connection of individual duplex ports to support the routing of specific switch ports to devices that fall outside of the MPO trunk groups.

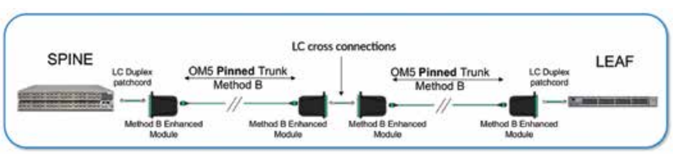

Figure 13: MPO 24 LC duplex

While two additional LC-LC connections are added to the channel, the increased flexibility can be important for your specific network design. The panel designs support duplex and parallel configurations as the panels are sized for duplex LCs.

Replacing MPO-duplex modules with array cords as illustrated in figure 14 increases the panel connectivity density. For example, a twelve port duplex LC module can be replaced with an adaptor housing 8 MPO ports. With MPO-24 connectors/trunks each MPO port can support twenty-four fibers, for a total of 192 fibers in the same module space. Comparing this capacity to the 12-port duplex module, the module will now support ninety-six duplex ports—an 8:1 increase. The same module space would typically support three QSFP ports (using a 3 x MPO-8 module with two MPO-12 trunks) compared to twenty-four with MPO-24 trunks and MPO-8 array cords. Therefore, for QSFP applications MPO-24 supports a 3:1 increase in panel density when compared to MPO-8 ports, and an 8:1 increase in panel density when compared to MPO-12 conversion modules.

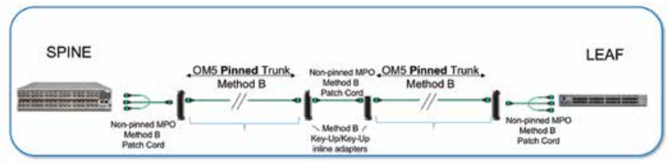

Figure 14: Fan-out arrays and assemblies increase cabling density

MPO-24 systems can also support an expanded range of parallel applications compared to MPO-8 or MPO-12 systems. The 100G SR-10 application requires ten pairs of MMF in a 10x10G configuration. Some manufacturers have extended this application to provide 12x10G switch ports. MPO-24 provides simple direct support for these 100G or 120G applications as illustrated in figure 15.

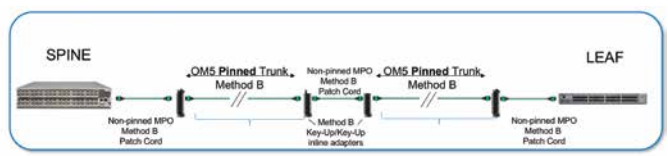

Figure 15: 100G and 120G parallel support with MPO-24

The 120G ports can be configured as individual 10G links—as server links as illustrated in figure 16. They can also be grouped into groups of four to support three 40G links. These applications are best supported with MPO-24 trunks and MPO-24 arrays or breakouts.

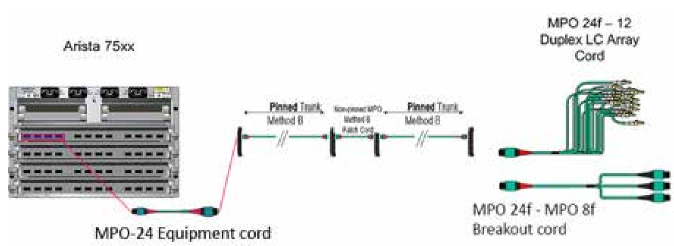

Figure 16: 120G parallel support with MPO-24

MPO system costs

We have reviewed the application support and configuration options in previous sections. The overall cost of MPO systems is influenced by the type of applications they will support. Flexible designs to accommodate either parallel or duplex applications require adequate rack space, cable routing and management as well as a combination of the correct media and performance levels. We have shown the design advantages of MPO-24 systems. Reducing the number of connections in the network also reduces the capital cost of the network as well as the installation and commissioning costs as illustrated in figure 17.

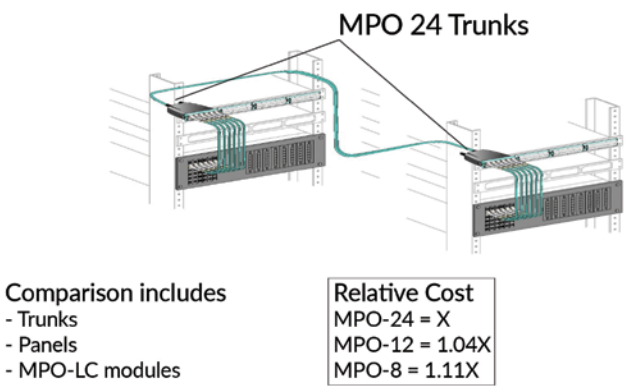

Figure 17: MPO system cost comparisons

Conclusions

Data center applications, networks and optical transceivers are evolving very quickly. This is in response to increasing demand for more services and efficiency to support exponentially increasing bandwidth requirements. These new network capabilities rely on the performance and adaptability of the physical fiber infrastructure. Optimizing the fiber infrastructure requires a high-speed migration strategy to support the emerging duplex and parallel optic transmission applications that will become popular as economic advantage shifts to newer technologies. The exact timing of these changes and the relative benefit to each data center will likely be somewhat different. Thus, the design objectives are speed, flexibility and scalable capacity.

MPO cabling systems are quick to deploy, flexible to configure, and scalable to optimal capacity. Manufactured with precision processes in factory settings, they provide excellent, reliable and repeatable application support for the most demanding network optics. MPO systems provide the solution to very challenging data center design requirements.

The basic MPO trunking systems are available in 8, 12 and 24-fiber variants. These MPO systems have evolved over time and now provide support for a large variety of network topologies. The large variation in network topologies has been driven by a constant shift in the economic options introduced by network equipment vendors. This trend continues with standards based applications and a variety of newly formed multi-source-agreement solutions all pushing data center technology forward. Exploring the potential applications that make sense for your data center, along with a MPO system that can support the widest array of options, is a sound strategy for maximizing the return on the fiber infrastructure investments you must make.

MPO SYSTEM CHARACTERISTICS

MPO-24

- The most efficient connector/trunk combination providing lower per port deployment costs
- High degree of flexibility for all duplex and parallel optical applications
- Lowers the number of connectors required in a system along with the cleaning and inspection costs associated
- Enables smaller cable diameters and reduces pathway fill requirements
- The superior choice for density with more ports per panel that any other MPO option

MPO-12

- The most familiar MPO connector/trunking cabling interface
- Lower density for duplex applications compared to MPO-24 systems
- Better density for duplex applications compared to MPO-8 systems

MPO-8

- Efficient for point to point QSFP trunking applications
- Convenient for QSFP breakouts in some specific applications
- Least efficient for duplex applications
- Offers the lowest panel density of all MPO systems

CommScope (NASDAQ: COMM) helps design, build and manage wired and wireless networks around the world. As a communications infrastructure leader, we shape the always-on networks of tomorrow. For more than 40 years, our global team of greater than 20,000 employees, innovators and technologists has empowered customers in all regions of the world to anticipate what's next and push the boundaries of what's possible.

Discover more at commscope.com

COMMSCOPE®

commscope.com

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

AN-111782.1-EN (10/17)