White paper

Data center High Speed Migration: Infrastructure issues, trends, drivers and recommendations

May 2017

Contents

Introduction

In the data center, speed is everything. The challenge is to look ahead and know what you have to be prepared to deliver—in the immediate future and later on—and chart the most expedient and flexible course forward. The faster that available technologies and applicable standards evolve, the harder that job becomes.

Recent data-center trends predict bandwidth requirements will continue growing 25 percent to 35 percent per year. A key impact of this sustained growth is the shift to higher switching speeds. According to a recent study by Dell'Oro, Ethernet switch revenue will continue to grow through the end of the decade, with the biggest sales forecasted for 25G and 100G ports.[1]

40G port revenue has likely peaked and will now decline in favor of 25G and 50G ports. 40G ports are used extensively for server connectivity today, with each QSFP 40G port supporting four 10G server connections. Servers are quickly evolving, however, and new designs will easily consume much more than a 10G uplink can supply.

The shift to 25G lanes is well underway as switches deploying 25G lanes become more commonplace. Lane capacities are expected to continue doubling, reaching 100G by 2020 and enabling the next generation of high speed links for fabric switches. A number of factors are driving the surge in data center throughput speeds.

- Server densities are increasing by approximately 20 percent a year.
- Processor capabilities are growing, with Intel recently announcing a 22-core processor.
- Virtualization density is increasing by 30 percent[2], which is driving the uplink speeds to switches.
- East-west traffic in the data center has far surpassed the volume of north-south traffic.[3]

"The idea going forward is to push lanes to 25 Gb/sec speeds, as the current crop of Ethernet switches are doing, and then ramp up to 50 Gb/ sec lanes and then 100 Gb/ sec lanes and keep the lane count down around eight."

– The Next Platform, March 2016

The network design has to reflect this massive amount of traffic, and, importantly, has to allow for server, storage, and network capacity to all be scaled up independently and with as little disruption and reconfiguration as possible. As a result, data center professionals must support higher server densities, deploy more fiber and accelerate plans to migrate to higher speeds in their core and aggregation networks. The network infrastructure within the data center must be able to scale to support these significant changes.

Changing network architecture

The change in data center traffic and direction require a network design that accommodates the rapid increase of east-west data traffic.

"Adoption of network architectures such as spine and leaf… are driving not only bandwidth demand, but also the scale of the network, requiring a greater fiber count for the cabling infrastructure."

– Data Center Journal, April 25, 2016

The traditional data center architecture used a three-layer topology as shown in Figure 2. The core layer, typically located in the main distribution area (MDA), is where the various network switches connect to each other, as well as to network sources outside the data center. This layer feeds the aggregation layer, connecting the various access switches. In large enterprise and cloud data centers, the aggregation layer is usually located in the intermediate distribution area (IDA). In smaller facilities it is typically the horizontal distribution area (HDA) or equipment distribution area (EDA). The access layer runs from the data center to the individual nodes where users connect to the network.

The design of this model provides a predictable foundation for a scalable data center network but is less than ideal when it comes to supporting today's low-latency, virtualized applications. As a result, there has been a swift and dramatic shift to the "leaf-and-spine" architecture.

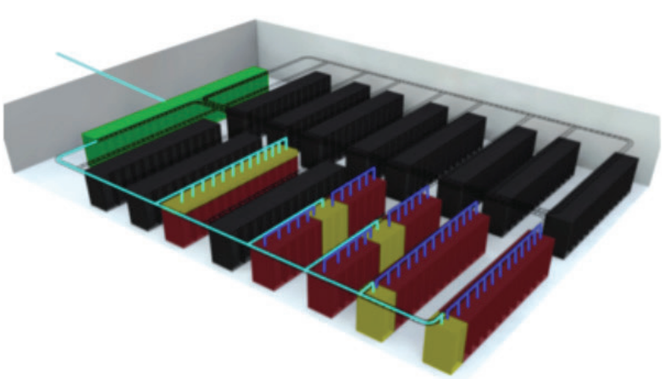

Figure 1: Data Center Design

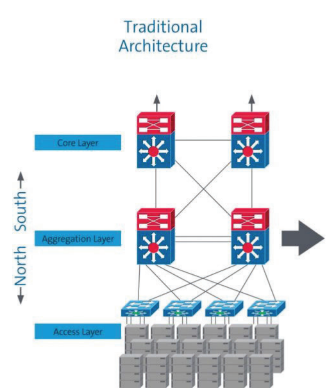

Figure 2: Traditional three-layer topology

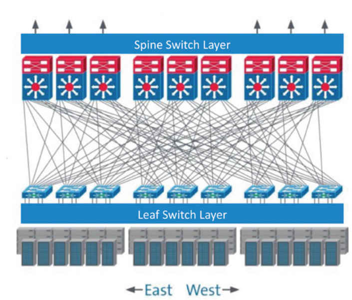

Figure 3: Leaf-and-spine topology

Illustrated in Figure 3, the leaf and spine model is a streamlined network design that moves data in an east-west flow, enabling servers to co-operate in delivering cloud-based applications. In this topology, networks are spread across multiple leaf switches, making the leaf-and-spine switch layer critical for delivering maximum scale and performance.

Each leaf switch is connected to every spine switch, creating a highly resilient any-to-any structure. The mesh of fiber links creates a high-capacity network resource or "fabric" that is shared with all attached devices. All fabric connections run at the same speed. The higher the speed, the higher the capacity of the fabric.

Fabric networks require a large number of fiber connections, particularly in the switch layer. Equipment vendors continuously work to increase the density of their line cards in order to keep pace. Given the increasing density, cabling connectivity and management solutions such as optical distribution frames, panels and raceways become more important.

Evolving standards

Applications standards organizations, namely the IEEE 802.3 (Ethernet) and ANSI/T11 (Fibre Channel Committees) have been busy updating the recommended guidelines in order to keep pace with the rapid bandwidth increases. The objective of these standards groups is not just to facilitate the evolution to ever-increasing line rates. They also encourage development of higher-speed applications that will increase the cost effectiveness of links between the data center equipment. To this end, a number of intermediate speeds are being developed to fill the gap between 10G, 40G, 100G and 400G. Table 1 lists the various Ethernet standards. Those still in process are listed in cyan.

TABLE 1— IEEE 802.3 ETHERNET FIBER STANDARDS—COMPLETED AND IN PROGRESS (CYAN)

Application	Standard	IEEE reference	Media	Speed	Target distance
10-Gigabit Ethernet	10GBASE-SR	802.3ae	MMF	10 Gb/s	33 m (OM1) to 550 m(OM4)
	10GBASE-LR		SMF		10 km
	10GBASE-LX4		MMF		300 m
	10GBASE-ER		SMF		40 km
	10GBASE-LRM	802.3aq	MMF		220 m (OM1/OM2) to 300 m (OM3)
25-Gigabit Ethernet	25GBASE-SR	P802.3by	MMF	25 Gb/s	70 m (OM3) 100 m (OM4)
40-Gigabit Ethernet	40GBASE-SR4	802.3bm	MMF	40 Gb/s	100 m (OM3) 150 m (OM4)
	40GBASE-LR4		SMF		10 km
	40GBASE-FR		SMF		2 km
	40GBASE-ER4		SMF		40 km
100-Gigabit Ethernet	100GBASE-SR10		MMF	100 Gb/s	100 m (OM3) 150 m (OM4)
	100GBASE-LR4		SMF		10 km
	100GBASE-SR4		SMF		70 m (OM3) 100 m (OM4)
	100GBASE-ER4		SMF		40 km
50G, 100G and 200G	50GBASE-SR	802.3cd	MMF	50 Gb/s	100 m (OM4)
	Ethernet		SMF		2 km
	50GBASE-LR		SMF		10 km
	100GBASE-SR2		MMF	100 Gb/s	100 m (OM4)
	100GBASE-DR2		SMF		500 m
	100GBASE-FR2		SMF		2 km
	200GBASE-SR4		MMF	200 Gb/s	100 m (OM4)
200-Gigabit Ethernet	200GBASE-DR4	P802.3bs	SMF		500 m
	200GBASE-FR4		SMF		2 km
	200GBASE-LR4		SMF		10 km
400-Gigabit Ethernet	400GBASE-SR16		MMF	400 Gb/s	70 m (OM3) 100 m (OM4)
	400GBASE-DR4		SMF		500 m
	400GBASE-FR8		SMF		2 km
	400GBASE-LR8		SMF		10 km

Options for migration

The discussion surrounding migration to higher line rates is both complex and rapidly evolving. It includes a wide range of decisions regarding fiber type, modulation and transmission schemes, connector configurations and, of course, cost considerations. Figures 4 and 5 show two possible migration paths, but there are many others. Determining which one is best for any given environment means carefully considering each aspect.

The following are just a few of the many issues that must be weighed.

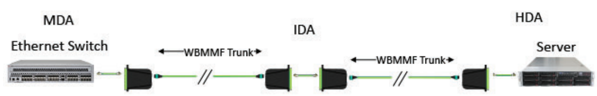

Figure 4: 10GBASE-SR duplex fiber link with MPO trunks

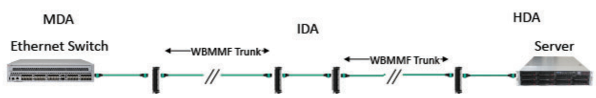

Figure 5: 40GBASE-SR4 link with parallel optics in switch and server

40G or 25G lanes?

Until recently, the accepted migration road map outlined a predicted jump from 10G lanes to 40G. Since the approval of the IEEE 802.3by standard, the industry has shifted to 25G lanes as the next switching technology. This is largely due to the fact that the newer 25G lanes, which offer easy migration to 50G (2x25G) and 100G (4x25G), provide a better return on investment than migrating to 40G.

Modulation schemes

New, more efficient modulation schemes are also now available. Pulse-amplitude modulation with four amplitude levels (PAM-4) has been proposed for either intra- or inter-data center optical links. As shown in Figure 6, PAM-4 is a modulation technique that uses four distinct pulse amplitudes to transmit data. Compared to traditional NRZ (binary modulation), PAM-4 enables twice the transmission capacity for the same signaling rate. The downside, however, is that it requires a higher signal-to-noise ratio (SNR), which puts stricter requirements on the supporting physical infrastructure. Still, its simplicity and low power consumption make PAM-4 modulation one of the most promising modulation techniques for 100G and beyond.

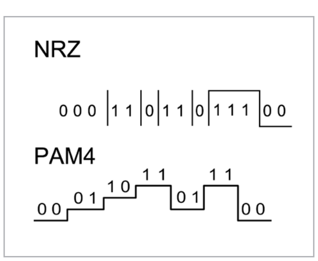

Figure 6: PAM4 and NRZ modulation

Transceiver technology

In addition to developing more advanced modulation schemes to increase various channel speeds, various wavelength division multiplexing (WDM) techniques are being developed to increase the number of lanes, or wavelengths transmitted on each fiber. WDM has been used for over two decades to increase data rates on long-haul networks by reducing fiber counts. It has also been used in singlemode Ethernet applications, such as 10GBASE-LR4 and 100GBASE-LR4, which combine four wavelengths on the same fiber using coarse WDM technology. This concept has also been extended to multimode fiber through what is known as shortwave WDM or SWDM. As shown in Figure 7, SWDM utilizes wavelengths from 850 nm to 940 nm.

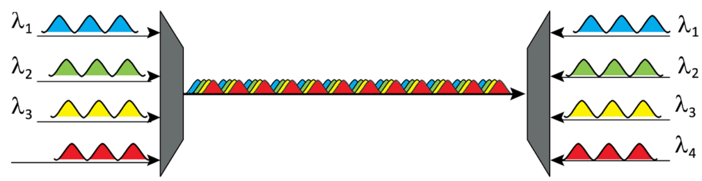

Figure 7: SWDM combining four wavelengths from 850 nm to 940 nm

Serial or parallel transmission?

As more demanding applications drive data rates higher, the market is gravitating to parallel optics. This trend is supported by the consistent demand for MPO-based trunks, a data center staple for more than a decade. Using laser-optimized multimode fiber (LOMMF), serial optics can cost-effectively support speeds up to 10G. But, historically, using serial transmission to support migration to 25G or 40G required switching to costlier singlemode solutions. Parallel optics, however, does provide a more cost-effective solution for migrating to 40G, 100G and 200/400G Ethernet.

The switch to parallel optics is enabled by the use of MPO connectors. In North America, sales of MPO fiber-optic connectors for 40/100GbE communication links are forecast to increase 15.9 percent annually through 2020, reaching $126 million in 2020.[4] However, the trend to parallel optics may ebb and flow as new technologies are implemented that make better use of duplex pairs. Technologies such as PAM4 and SWDM are expected to provide more cost-effective support for duplex data center applications in the near future.

Preterminated vs field-terminated cables?

The need to turn up networking services quickly has increased the value and demand for preterminated cabling systems. By some estimates, the plug-and-play capability of preterminated cables translates to 90 percent time savings versus a field-terminated system and are about 50 percent faster when it comes to network maintenance.[5] The value grows as the number of fiber connections within the network increases.

Factory-terminated systems are also the only viable solution to the extremely low-loss systems that are required to support high speed optic links. Among preterminated solutions, MPO/MTP fiber is fast becoming the de-facto system for both single- and multimode connectivity due to its ease of use and speed, not to mention the high density.

Singlemode, multimode or wideband multimode?

The cost of the pluggable optics continues to limit the implementation of singlemode fiber (SMF) in data centers. Although new technologies and manufacturing efficiencies are helping to reduce prices for SMF, it is still not enough to justify the high cost of singlemode optics. Two areas of the data center where use of SMF is increasing are from the entrance facility to the main distribution area and for extreme scale in mega-data center designs.

For the enterprise data center, multimode fiber (MMF) continues to offer a more attractive balance of performance, density and cost. The challenge for MMF is distance. As data traffic grows and interconnectivity speeds increase, the maximum distance for a communication link tends to decrease. But emerging higher quality components and engineered links can provide the link capacity to support the longer distances and new data center topologies.

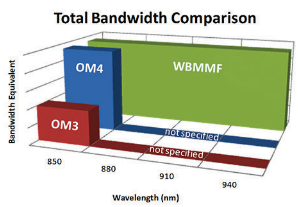

Figure 8: Fiber bandwidth comparison

The recent introduction of OM5 may eventually provide the optimum solution for fiber migration. Introduced by CommScope in 2015, OM5 was recently approved under ANSI/TIA-492AAAE and is expected to be recommended by ANSI/TIA-942-B. The new fiber enhances the ability of short-wavelength division multiplexing (SWDM) technology to provide at least a four-fold increase in usable bandwidth. It also supports all legacy multimode applications by maintaining compatibility with OM3 and OM4 fiber. By multiplexing four wavelengths spaced in the 850–950 nm region, one strand of WBMMF can increase the data capacity by a factor of four. Figure 8 shows a bandwidth comparison between OM3, OM4 and OM4 wideband fibers.

Intelligent systems

Automated infrastructure management (AIM) systems can greatly assist in the migration process by providing an accurate mapping of the physical layer and all connected devices. Because AIM systems automatically monitor and document all ports and fibers in use, they can help ensure capacity is available when upgrading from duplex to parallel.

Conversely, AIM can help identify surplus cabling and switch ports and make them available for parallel-to-duplex migration. The ISO/IEC 18598 standard as well as the European Standard EN 50667 for AIM were ratified in 2016, and the ISO/IEC document is being adopted nearly verbatim by TIA as ANSI/TIA-5048.

The CommScope view

Preterminated MPO-based fiber solutions continue to be the best choice for high- performance networks. These systems provide excellent factory-terminated performance, plus the speed and agility to support the expansion requirements of private, cloud-like enterprise data centers.

Increasing the bandwidth from OM3 to OM4 is another "must have" for today's high-capacity networks. Of the MMF options, OM5 is the ultimate selection due to its ability to increase the practical capacity of OM4 by a factor of four. Also, in order to increase the effective density of networking equipment and the supporting physical network we support the use of SWDM duplex for fabric links.

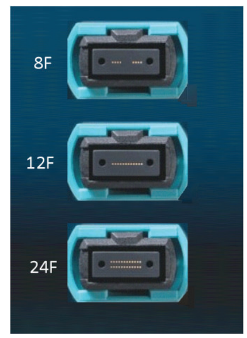

Figure 9: MPO connectors with varying fiber counts

There also may be a bit more clarity in the on-going debate between eight-fiber, 12-fiber and 24-fiber MPO technology. MPO 12-fiber systems have been deployed for years. They support duplex and parallel applications with superior flexibility and reach for most data center applications. The benefit of operational uniformity would argue for a continued use of this same system for future applications. Certainly increasing the bandwidth from OM3 to OM4 and ultimately to OM5 is a benefit for higher capacity networks in the future.

MPO 24-fiber systems, meanwhile, increase the density and capacity of physical networks. These systems support duplex and parallel applications and offer lower costs per fiber compared to eight-fiber and 12-fiber systems. As a result, this is the recommended system for high-density networks or networks with primarily duplex applications.

MPO eight-fiber systems support popular four-lane QSFP applications used predominantly in 4X10G or 4X25G configurations for storage and server network attachments.

Network fabric links do not require breakouts to lower speed ports. Therefore, two-fiber duplex links such as 100G SWDM are a very attractive choice. These duplex fiber fabric links are implemented with higher density, using either 12-fiber or 24-fiber solutions.

With the MPO 24-fiber system as the preferred implementation, CommScope's solutions support eight-fiber parallel and two-fiber duplex applications offering the optimal support for a broad array of data center applications. MPO 24 fiber systems provide optimized support for both duplex and parallel applications and can provide lower overall costs while supporting the variety of up-coming optic applications.

A closing thought

While it is important to understand the vast array of technical options and emerging solutions, these must be viewed within the context of your specific enterprise data center environment. What is the trajectory of the enterprise? How does that affect the velocity of change and scaling requirements in the data center? What is the total cost of ownership for the various migration scenarios being considered?

As a data center manager, remember, you don't have to go it alone. The amount of research and decisions involved can be mind-numbing. There are a variety of knowledgeable resources, such as CommScope, who have the solutions and experience to help you make the right decision. By leveraging our technical expertise and broad perspective, together we can help you develop a long-term migration strategy designed to keep your data center adaptable, capable and efficient. No matter how fast things change.

Sources

[1] Construction Zones on the Ethernet Roadmap; The Next Platform; March 24, 2016

[2] Data Center Strategies North American Enterprise Survey; Infonetics Research; May 2015

[3] Facebook Gives Lessons In Network-Datacenter Design; November 2014

[4] Market Forecast—MPO Connectors in 40/100GbE; ElectroniCast Consultants; December 2015

[5] Weighing the costs and benefits of preterminated fiber-optic systems; Cabling Installation & Maintenance; May 1, 2014

[6] PAM-4: A Key Solution for Next-Generation Short-Haul Optical Fiber Links; Neophotonics, blog; May 11, 2015

Everyone communicates. It's the essence of the human experience. *How* we communicate is evolving. Technology is reshaping the way we live, learn and thrive. The epicenter of this transformation is the network—our passion. Our experts are rethinking the purpose, role and usage of networks to help our customers increase bandwidth, expand capacity, enhance efficiency, speed deployment and simplify migration. From remote cell sites to massive sports arenas, from busy airports to state-of-the-art data centers—we provide the essential expertise and vital infrastructure your business needs to succeed. The world's most advanced networks rely on CommScope connectivity.

commscope.com

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

WP-110615.1-EN (05/17)