

Six Considerations When Planning Data Center Network Upgrades

by Robert Reid from Panduit
and Pat Chou from Cisco



As applications like AI, 5G, and IoT continue to grow, so will the need for processing power and bandwidth in the data centers that support their associated applications. It's no surprise that these data centers are rapidly migrating from architectures based on 10 Gb data rates per fiber optic link, to ones based on 40 Gb and 100 Gb.

Many data center operators are under pressure to ensure that their network hardware and fiber cable infrastructure strategy not only supports today's requirements, but also provides a cost-effective upgrade path to accommodate the inevitable future growth. The key to minimizing total long-term cost is to consider both the fiber cable infrastructure and optical transceivers together. With wise planning, especially prior to upgrading to 40 Gb or 100 Gb, data center operators can avoid costly maintenance and changes in fiber cable infrastructure for years to come.




The first decision in a prudent strategy is to use BiDi (bi-directional) pluggable optical transceivers at 40 Gb or 100 Gb. The second is to use an enhanced OM4 MMF (multi-mode fiber) instead of standard OM4 MMF. If you're a data center operator, you should consider the benefits of this approach. They're captured by the following six considerations:

1. Greenfield installation: Choosing BiDi reduces fiber cable infrastructure cost of new deployments

The most common architecture in the modern data center is leaf-spine. With only two tiers, it ensures maximum two-hop server-to-server connectivity, which is important for low-latency application-to-application communication, virtualized LANs, and virtualized servers. Because leaf-spine requires many links to connect all leaf switches directly to all spine switches, economizing these links has a significant impact on reducing total cost.

Consider a spine layer composed of two Cisco Nexus 9364C 64-port switches, and a leaf layer consisting of 32 Cisco Nexus 93180YC switches (For illustration purposes, only four are shown in Figure 1), each with six 100 Gb uplink ports available. You could support 2 links between each leaf switch and each spine switch, for a total of 128 links at 100 Gb. With a chassis-based modular switch, even more links per leaf-spine combination could be supported for even higher throughput.



Cisco Nexus 9000 Ethernet switches in leaf-spine architecture.

With this many links between leaf and spine, fiber cable infrastructure to support SR4 transceivers would cost substantially more than the infrastructure for BiDi by about 80%. This is attributed to the substantial number of trunks and jumpers required to support SR4, and because MPO connectors are more costly than LC connectors.

2. Brownfield installation: Upgrading to BiDi transceivers avoids the need to modify fiber cable infrastructure

Ideally, a data center network upgrade process should involve replacing only the servers, ethernet switches, and pluggable transceivers. However, transceivers based on the IEEE standardized 40 and 100 Gb SR4 optical interface specifications use parallel fiber and MPO connectors. 10 Gb SR uses serial duplex fiber and dual-LC connectors. Therefore, upgrading from 10 Gb SR to 40 or 100 Gb SR4 would require replacing jumpers, adding more trunk cabling, and adding MPO-based connectivity at the transceivers. This is illustrated in Figure 2 below.

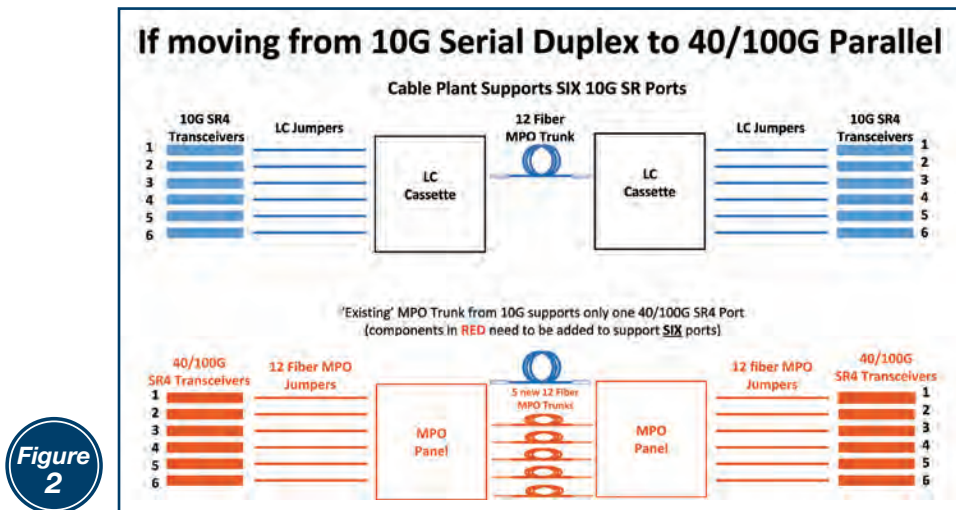


Figure 2

Upgrading to 40 or 100 Gb links with IEEE standardized SR4 pluggable transceivers. Dual-LC jumpers need to be replaced with MPO jumpers, and more fiber needs to be added to the trunks.

In contrast, 40 and 100 Gb BiDi transceivers allow you to leverage the same LC-based duplex MMF infrastructure used with 10 Gb SR transceivers. This means you can upgrade the network gear without making changes to the fiber cable infrastructure. This is illustrated in Figure 3.

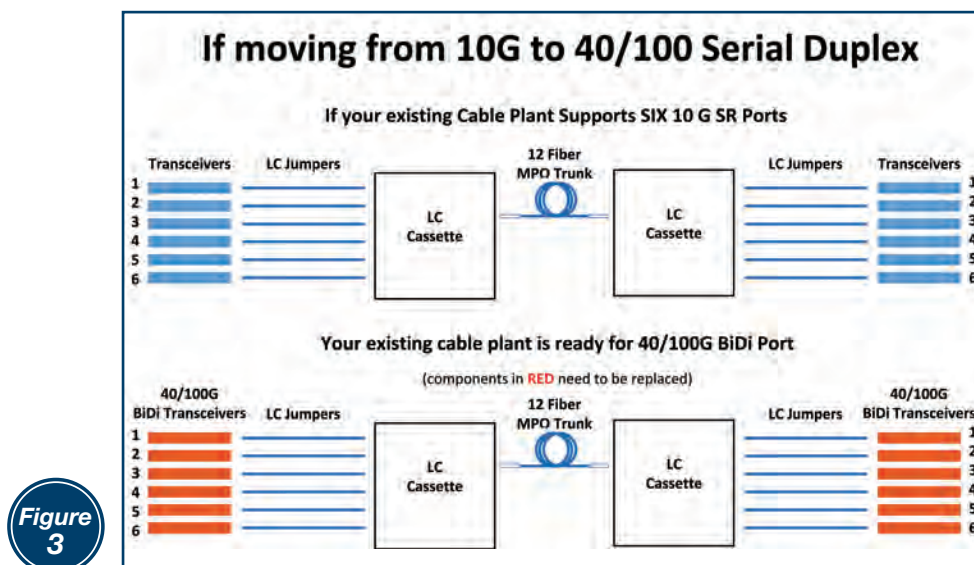




Figure 3


Upgrading to 40 or 100 Gb links with BiDi pluggable transceivers. No change to the LC jumpers or MPO trunk cabling.



MPO connector end face. This connector is used for transceivers that use four fiber pairs (eight fiber strands) to transmit four channels in parallel, such as SR4. Four of the 12 fibers are used for the transmit direction, and four for the receive direction. The four fibers in the middle are unused by the transceiver.



Dual-LC connector end faces. This connector is used by 10G SR transceivers, which transmit 10 Gb data in both directions over duplex fiber cable. Each of the two fibers carries one direction.



Dual-LC connector end faces. This is the same connector as in Figure 5 and is also used by 40G and 100G BiDi transceivers. Unlike 10G SR, the BiDi transceivers transmit in both directions on each fiber. The 40 Gb BiDi transmits 20 Gb in each direction on both fibers, for a total of 40 Gb data rate. Similarly, the 100 Gb BiDi transmits 50 Gb for a total of 100 Gb data rate.

3. Operational costs: Duplex fiber and LC connectors keep it simple

Fiber cable infrastructure includes not only the fiber optic cable, but also fiber optic connectors. If you choose duplex fiber infrastructure over parallel fiber infrastructure, you can avoid using MPO connectors to connect to the transceivers. Benefits of LC over MPO include:

- Easier to clean and inspect
- Less variability in optical insertion loss
- Less likely to be contaminated
- MACs (moves, adds, and changes) are easier

These translate to lower maintenance costs.



4. Switch-to-switch fiber lengths: **Enhanced OM4 + BiDi should satisfy your link distances**



Consider the size of your data center. Unlike hyperscale data centers, which have large footprints and often use more expensive single-mode fiber transceivers for longer reach, the rest of the world's data centers aren't nearly as big and seldom have links longer than 100m. In fact, only 10% of all data center links are longer than 100m, according to the Ethernet Alliance.

So, what to do about that 10%? 10 Gb SR transceivers reach well beyond 100m. At 40 Gb, SR4 and BiDi reach is 150m over standard OM4 MMF. But at 100 Gb, the reach over standard OM4 MMF is 100m, and only 70m over standard OM3 MMF. If you use an enhanced OM4 MMF that is optimized for BiDi wavelengths, you can easily exceed well over 100m even with a 100 Gb BiDi transceiver.

5. Cable infrastructure design: **LC fiber connection points maximize link distance and flexibility**




At fiber connection points, for instance at patch panels, optical fiber interfaces in LC connectors have lower optical signal loss than those of MPO connectors, which are needed by SR4 transceivers. In cases where a link contains so many connection points that the loss budget is exceeded, the link is still feasible at a distance shorter than the reach. Therefore, in this situation, LC connectors enable longer distance than MPO for a given number of connection points because LC connection points lose less optical signal. Alternatively, LC connectors support more connection points over a given distance, which means that more patch panels can be used.

6. Fiber cable infrastructure: **Re-use your duplex MMF LC based infrastructure through 100G and beyond**



In general, due to the physics of the light pulses traveling through MMF, reach tends to decrease with increasing data rates. This means that a link that was short enough for 10 Gb SR may be too long when you upgrade to 100 Gb. To future-proof your fiber cable infrastructure, install an enhanced OM4 fiber that is optimized for BiDi wavelengths. If you install this when you start using 10 Gb, you can leave it untouched for years as you upgrade to 100 Gb BiDi, and potentially through 400 Gb.



Leading the Pack in Upgrades to 100 Gb and beyond

One organization that has been successful in upgrading its data centers is the Hertz Corporation. Hertz needed to upgrade its data center from 10 Gb to 100 Gb to accommodate expected growth while maximizing ROI, reducing operating expenses, and integrating sustainability throughout the business infrastructure in line with the company's corporate efficiency goals. The upgrade included a design of all pathways and cabinet layout in a new white space as well as a migration pathway that would move seamlessly from the existing space to the new space. Partnering with Panduit and Cisco allowed for Hertz to achieve its network infrastructure implementation with sustainable operations and added scalability for future growth.

Other ambitious programs have demonstrated a clear ROI. For example, a major U.S. financial institution with a global footprint was looking for a cost-effective data center solution to migrate to a 100 Gb data rate for high-speed trading. Working with Panduit to align its optics and fiber infrastructure, this organization achieved manageability and the lowest latency possible in its network. The organization became one of the first in the industry to reach 100 Gb, delivering the fastest speeds and reliability in the financial services industry, giving it an edge over its competitors.

Summary

Data center operators can avoid enormous cost and complication over several years if they have installed an enhanced OM4 MMF cable plant in their fiber infrastructure, and plan to upgrade to 40 or 100 Gb using BiDi optical transceivers in their ethernet switches. The cost savings are realized in material spend, installation labor, and maintenance cost, thanks to BiDi transceivers and enhanced OM4 MMF enabling flexible design options and connectivity that is easy to service. Thus, the same fiber cable infrastructure can remain untouched used as operators migrate to 40 Gb, 100 Gb, and possibly to 400 Gb.

Panduit Signature Core MMF is optimized for Cisco BiDi transceivers and supports greater reach than standardized OM4 MMF. Cisco offers the industry's only dual-rate 40/100 Gb BiDi transceivers that enables flexibility and incremental upgrades to 100 Gb from 40 Gb.



PANDUIT®

Since 1955, Panduit's culture of curiosity and passion for problem solving have enabled more meaningful connections between companies' business goals and their marketplace success. Panduit creates leading-edge physical, electrical, and network infrastructure solutions for enterprise-wide environments, from the data center to the telecom room, from the desktop to the plant floor. Headquartered in Tinley Park, IL, USA and operating in 112 global locations, Panduit's proven reputation for quality and technology leadership, coupled with a robust partner ecosystem, help support, sustain, and empower business growth in a connected world.

For more information

Visit us at www.panduit.com

Contact Panduit North America Customer Service by email: cs@panduit.com
or by phone: 800.777.3300

THE INFORMATION CONTAINED IN THIS ARTICLE IS INTENDED AS A GUIDE FOR USE BY PERSONS HAVING TECHNICAL SKILL AT THEIR OWN DISCRETION AND RISK. BEFORE USING ANY PANDUIT PRODUCT, THE BUYER MUST DETERMINE THE SUITABILITY OF THE PRODUCT FOR HIS/HER INTENDED USE AND BUYER ASSUMES ALL RISK AND LIABILITY WHATSOEVER IN CONNECTION THEREWITH. PANDUIT DISCLAIMS ANY LIABILITY ARISING FROM ANY INFORMATION CONTAINED HEREIN OR FOR ABSENCE OF THE SAME.

All Panduit products are subject to the terms, conditions, and limitations of its then current Limited Product Warranty, which can be found at www.panduit.com/warranty.

*All trademarks, service marks, trade names, product names, and logos appearing in this document are the property of their respective owners.

PANDUIT US/CANADA
Phone: 800.777.3300

PANDUIT EUROPE LTD.
London, UK
cs-emea@panduit.com
Phone: 44.20.8601.7200

PANDUIT SINGAPORE PTE. LTD.
Republic of Singapore
cs-ap@panduit.com
Phone: 65.6305.7575

PANDUIT JAPAN
Tokyo, Japan
cs-japan@panduit.com
Phone: 81.3.6863.6000

PANDUIT LATIN AMERICA
Guadalajara, Mexico
cs-la@panduit.com
Phone: 52.33.3777.6000

PANDUIT AUSTRALIA PTY. LTD.
Victoria, Australia
cs-aus@panduit.com
Phone: 61.3.9794.9020