Making Strategic Decisions About Optics and FC Cabling for System z Enterprises

By David Lytle

What optics and cabling are best used in data center environments? This topic remains a primary concern for many users. And since there's a cost associated with optics and cabling deployment, this really is a serious decision that shouldn't be made lightly. Figure 1 shows that each generation of Fibre Channel adds significantly more ports into data centers which, at this point, require serious attention to the types and structure of optics and cables.

Industry groups have done a great deal of work regarding data center cabling. For instance, the ISO/IEC 24764 (2010) Information Technology—Generic Cabling Systems for Data Centers International Standard, in conjunction with ISO/IEC 11801 Edition 2.2, specifies both balanced copper and optical fiber cabling systems targeted for deployment in the data center. This standard specifies a modified structure and configuration for generic cabling within data centers, as well as a reference implementation specific to data center infrastructures. This standard provides data center users with a flexible and generic cabling system that's easy to modify and capable of supporting a wide range of applications.

But the question we often hear from System z users, and sometimes open systems users, is more specific in nature: What are others doing as far as deploying longwave (LX) and/or shortwave (SX) optics and their associated single mode (SM) or multimode (MM) cables? To better understand the answer to that question, let's briefly discuss the FICON channel features, switch ports and storage adapters. Note that optics, as used in this discussion, are small form-factor pluggable (SFP) or SFP+ connections.

FICON Channel Features, Switch Ports and Storage Adapters

Figure 2 succinctly illustrates the point that each generation of Fibre Channel has not only advanced its optical tempo but also generated new requirements for optical transceivers that take advantage of faster connection speeds. IBM zEnterprise product managers obviously match the capability of their I/O system with one or more of these Fibre Channel generations. Today, IBM's FICON Express8 and FICON Express8S channel features, supporting up to 8Gbps, come with either all SX or all LX optics, since SX and LX optics can't be intermixed on the same FICON feature card. The cost for either card is identical, but LX ports allow direct connectivity up 10 kilometers (km) from where the mainframe sits. So, LX optics provide more of a benefit than the local-distance SX optics. With this in mind, many mainframe shops worldwide will deploy FICON Express8 or FICON Express8S feature cards with LX optics.

However, LX switch and storage optics for ports do cost more than their SX cousins. This additional cost is typically the focus of the conversation. Small enterprises and open systems enterprises are often more focused on minimizing costs and simplifying management. The additional cost of LX optics and SM cabling for switches and storage can be difficult to justify. It requires a strategy that includes a policy of making long-term investment decisions that can overcome the immediate budgetary issues of higher-cost optics and cables.

When thinking about long-term benefits, it's necessary to recognize that Fibre Channel line rate speed is doubling about every two years. This means that over time, SX optics and MM cables become a constraint on data center configurations, since the supported cable distance for MM cables becomes increasingly shorter as line rate increases. For enterprises that opt for SX optics and MM cables, it might become necessary to replace existing MM cabling every four to six years to overcome the cable's distance constraints and realize all the value from new purchases of faster host and storage devices attached through that cabling.

SAN Generation	Release Date	# of FC Ports in a SAN	Square Footage	Speeds	Structured Cabling/ Multifloors
1st	1998	10s	1,000s	1GFC	No
2nd	2002	100s	10,000s	2GFC	Some
3rd	2005	1,000s	100,000s	4/10GFC	Common
4th	2008	10,000s with NPIV	100,000s*	8/10GFC	Very Common
5th	2011	100,000s with NPIV	100,000s*	16GFC	Required

Figure 1: Cabling the Next Generation SAN
Source: Brocade presentation

Generation	1st Gen	2nd Gen	3rd Gen	4th Gen	5th Gen
Electrical / Optical Module	1GFC / GBIC/ SFP	2GFC / SFP	4GFC / SFP	8GFC / SFP+	16GFC / SFP+
Electrical Speeds(Gbps)	1 lane at 1.0625	1 lane at 2.125	1 lane at 4.25	1 lane at 8.5	1 lane at 14.025
Encoding	8b/10b	8b/10b	8b/10b	8b/10b	64b/66b
Availability	1997	2001	2006	2008	2011
Connector	GBIC ←		SFP / SFP+		→

Figure 2: 16GFC Sets the Pace for Storage Networks
Source: Brocade presentation by Scott Kipp

	1 Gbps		2 Gbps		4 Gbps		8 Gbps		10 Gbps		16 Gbps	
Fiber Core (µ) Light source	Distance meters feet	* Link loss budget	Distance meters feet	* Link loss budget	Distance meters feet	* Link loss budget	Distance meters feet	* Link loss budget	Distance meters feet	* Link loss budget	Distance meters feet	* Link loss budget
9µ SM LX laser	10 km 6.2 miles	7.8 dB	10 km 6.2 miles	7.8 dB	10 km 6.2 miles	7.8 dB	10 km 6.2 miles	6.4 dB	10 km 6.2 miles	6.0 dB	10 km 6.2 miles	6.4 dB
9µ SM LX laser	4 km 2.5 miles	4.8 dB	4 km 2.5 miles	4.8 dB	4 km 2.5 miles	4.8 dB	N/A	N/A	N/A	N/A	N/A	N/A
50µ MM OM4 4700 MHz-km SX laser	As good as OM3	N/A	As good as OM3	N/A	400 m 1312 ft	2.95 dB	200 m 656 ft	2.19 dB	550 m 1805 ft	2.6 dB	130m 427 ft	1.95 dB
50µ MM OM3 2000 MHz-km SX laser	860 m 2822 ft	4.62 dB	500 m 1640 ft	3.31 dB	380 m 1247 ft	2.88 dB	150 m 492 ft	2.04 dB	300 m 984 ft	2.6 dB	100m 328 ft	1.86 dB
50µ MM OM2 500 MHz-km SX laser	500 m 1640 ft	3.85 dB	300 m 984 ft	2.62 dB	150 m 492 ft	2.06 dB	50 m 164 ft	1.68 dB	82 m 269 ft	2.3 dB	35m 115 ft	1.63 dB
62.5µ MM OM1 200 MHz-km SX laser	300 m 984 ft	3.00 dB	150 m 492 ft	2.10 dB	70 m 230 ft	1.78 dB	25 m 82 ft	1.58 dB	33 m 108 ft	2.4 dB	15m 49 ft	1.56 dB

Figure 3: FICON Implementation
Source: Fundamentals of Brocade Mainframe Networking Class

It should also be recognized that when using 8Gbps or faster optics, any time a cable is "touched," both ends of the cable need to be cleaned. When dealing with hundreds or thousands of new cables, this can take a substantial amount of time within the upgrade project to accomplish. Even brand new cables from the factory usually have dirty cable ends that must be cleaned prior to use to assure good connections can be made. To be clear, there's definitely a huge cost in terms of money and manpower when upgrading hundreds or thousands of cables and that's coupled with a very disruptive cable replacement that can impact system performance and availability. The old saying, "pay me now or pay me later," really applies in this decision-making process.

Figure 3 shows that as line rate increases, MM cable distance decreases. If, over time, a user went from 4Gb to 8Gb and then 16Gb, to keep the same distance between data source and the device, the MM cable would have to be replaced each time—or devices might need to be relocated.

Making an Investment for Long-Term Gain

Midsize to large mainframe enterprises are typically focused on a different set of issues than their smaller counterparts; high among them are usually predictability, performance and high availability. Predictability is important because an enterprise always wants to reap the value from whatever purchases they're making. If they buy a bigger host, it should provide equal or better application performance than before. If they buy new storage, it should provide equal or better application response time than before. So, buying MM cables, whose operational characteristics and predictability degrade over time as line rate speed increases, isn't feasible, since the enterprise wants its cabling to last as long as possible. Likewise, over time, a specific MM cable can only be 100 percent utilized if devices are within the supported distance or redeployed closer and closer to the data source as line rate speed increases. Not only is this potential relocation of equipment an unwelcome burden, it could affect the enterprise's overall system performance and lower or disrupt availability. SM cables don't require any equipment relocation, as they're designed to operate within a data center and up to 10km beyond. Many customers have been using their original SM cabling for more than 14 years.

With this in mind, LX optics and SM cables provide many more benefits over a longer period of time than the alternative. But the decision-maker must be willing to make an investment for long-term gain in order to justify the additional costs.

An Alternative … But With a Price

Of course, there's an alternative, a middle ground. It's possible, although not really desirable, to mix SX and LX SFP devices in the same FICON switch chassis. So, for example, if a customer wants to use LX optics on their System z CHPIDs but SX optics on their storage, that can be done. The FICON switch would need to have the correct number of LX and SX optics installed to support this. It would connect to the System z via a pair of LX optics and an SM cable and then attach to storage via a pair of SX optics using an MM cable. However, keeping track of and upgrading configurations over time becomes much more onerous for the user. And remember that SFPs are the most common failure within a storage area network. Users would need to have both LX and SX replacement optics on hand and be careful to ensure that the proper replacement optic is used whenever an SFP fails. Port swaps will also be more problematic, as they don't care about the optic type. For example, if a spare port on a switch is LX but a user accidentally swaps to it for a failed SX optic, the failing

> **Transceivers (SFPs) and fiber cables can be some of the most troublesome components of a storage network, so is there anything a proactive user can do, during deployment or troubleshooting activities, to easily diagnose bad SFPs and cables? Yes!**

port's MM cable will no longer match the swap port's LX optic. Intermixing LX and SX optics on the same switching device can become annoying.

Diagnosing Bad SFPs and Cables

Transceivers (SFPs) and fiber cables can be some of the most troublesome components of a storage network, so is there anything a proactive user can do, during deployment or troubleshooting activities, to easily diagnose bad SFPs and cables? Yes!

Diagnostic tools are available that ensure optical and signal integrity for Gen 5 Fibre Channel optics and cables, simplifying deployment and support of high-performance fabrics. These tools let users automate a battery of tests to measure and validate latency and distance across the switch links, as well as verify the integrity of the fiber cables and 10/16 Gbps SFPs in the fabric—either prior to deployment or when there are suspected physical layer issues.

Summary

Worldwide, budgets are tight, so a strategic investment in LX optics and SM cables could provide more benefits compared to the alternative. It isn't in any enterprises' best interest to spend budget money on an investment in cabling that probably has to be ripped out and replaced in only a few years' time. Numerous mainframe users have made the strategic decision to deploy all LX optics and SM cables in their enterprises. They realize that making this kind of an investment will minimize some significant risk and allow them to derive much more value within their I/O infrastructure in the long term compared to MM cables paired with SX optics. **ETJ**

David Lytle is a principal engineer/global solutions architect with Brocade Communications, Inc. in San Jose, CA. Based just outside of Atlanta, GA, he specializes in Gen 4 (8Gb) and Gen 5 (16Gb) Fibre Channel solutions and specifically in System z mainframe technologies, and provides sales enablement assistance to Brocade OEMs, partners and customers. David, a 45-year veteran of the IT industry, speaks at user conferences and is a co-creator and instructor for the Brocade Certified Architect for FICON (BCAF) professional certification class. He is a co-author of the recently published Brocade Bookshelf e-book *Brocade Mainframe Connectivity Solutions*, which is available at www.brocade.com/downloads/documents/books/brocade-mainframe-connectivity-bk.pdf.
Email: david.lytle@brocade.com