PANDUIT®
infrastructure for a connected world

Fibre Channel Rolls On

WHITE PAPER

For years, Fibre Channel has been marginalized as the industry's focus has turned to newer technologies such as Fibre Channel over Ethernet (FCoE) and IP-based storage networks. But Fibre Channel is very much alive and starting to grow again. The main catalyst for its continued use and relevance is the growth of cost-effective flash-based storage coupled with the availability of: 32 gigabit (Gb) Gen6 transceivers; 32 Gb Fibre Channel (GFC), 128 GFC and other technologies; and higher-capability multimode optical fiber cabling infrastructure.

For flash storage devices, the 32 Gb per second (Gb/s) line rate of Gen6 Fibre Channel is significant, as faster access and sustained read/write capability yield greatly improved transactional storage fabric throughput over previous generations of Fibre Channel.

This white paper covers current data center storage applications and technologies such as transceivers and optical fiber media that support Fibre Channel in the data center. It highlights the differences between Fibre Channel and competing technologies. The paper also discusses a roadmap for Fibre Channel (Gen 7 and beyond) in the context of future data center applications as well as the similarities and differences between the Fibre Channel and Ethernet higher speed roadmap and physical fiber cabling infrastructure needed to support both.

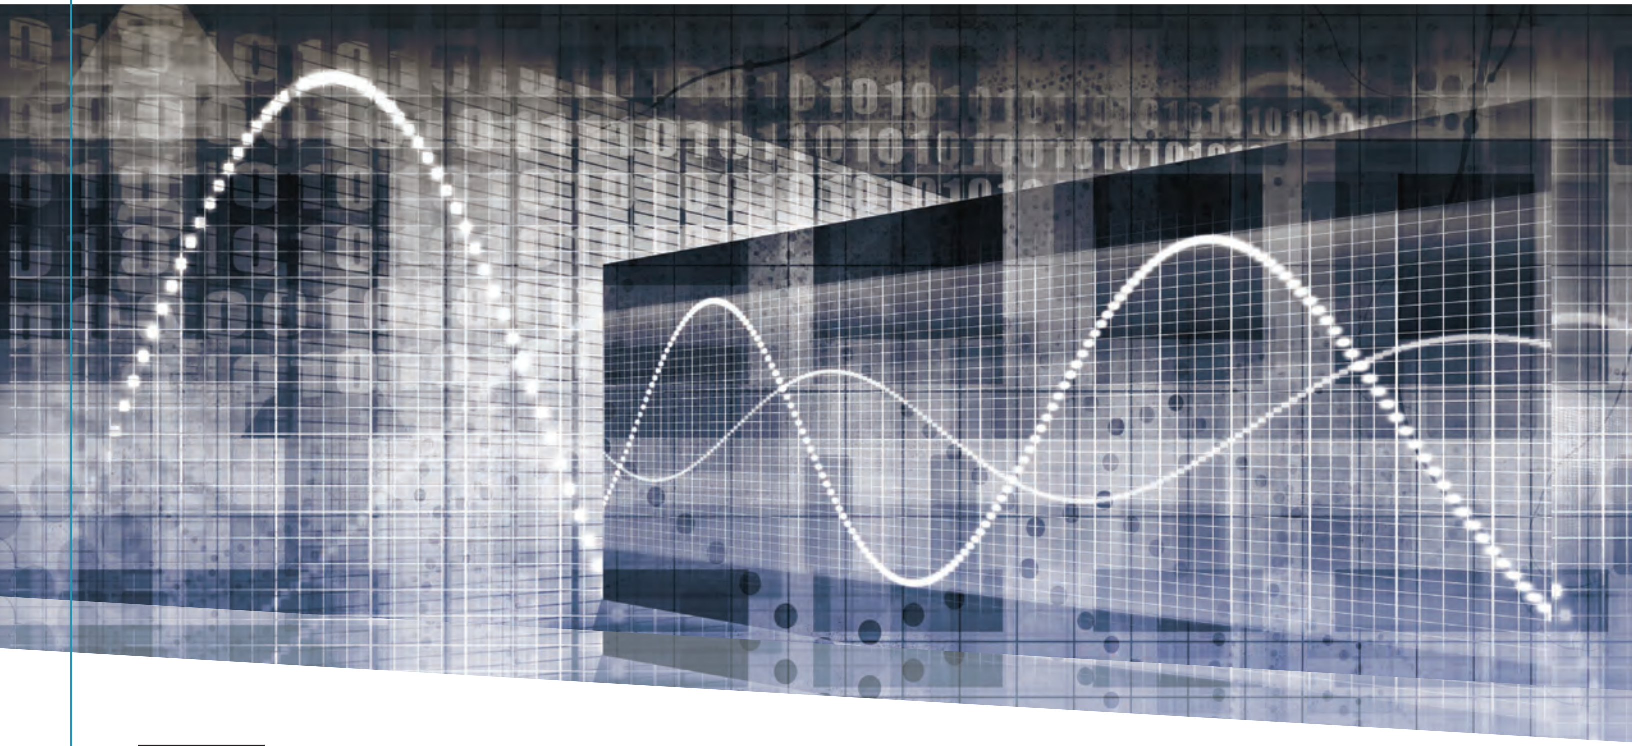

A Historical Perspective

The 1 Gb/s Fibre Channel standard was ratified in 1997 and new speeds have continued to evolve since (Figure 1).

The lifecycle of 1 GFC, from its market availability in 1998 until its market obsolescence in 2004, occurred simultaneously with the development and market release of 2 GFC and the first phase of standards development for 4 GFC. In 2005, transceivers supporting 2 GFC shipped in volume while 4 GFC was ratified as a standard and underwent general market release. This cycle of market acceptance, obsolescence, and renewal has continued over the last 20 years.

There is a significant concurrent overlap of data rates and technologies deployed, as the refresh cycle on customers' host and switching gear may be different (e.g., switches may have 16 GFC line cards but older storage arrays may only attach at 4 Gb/s).

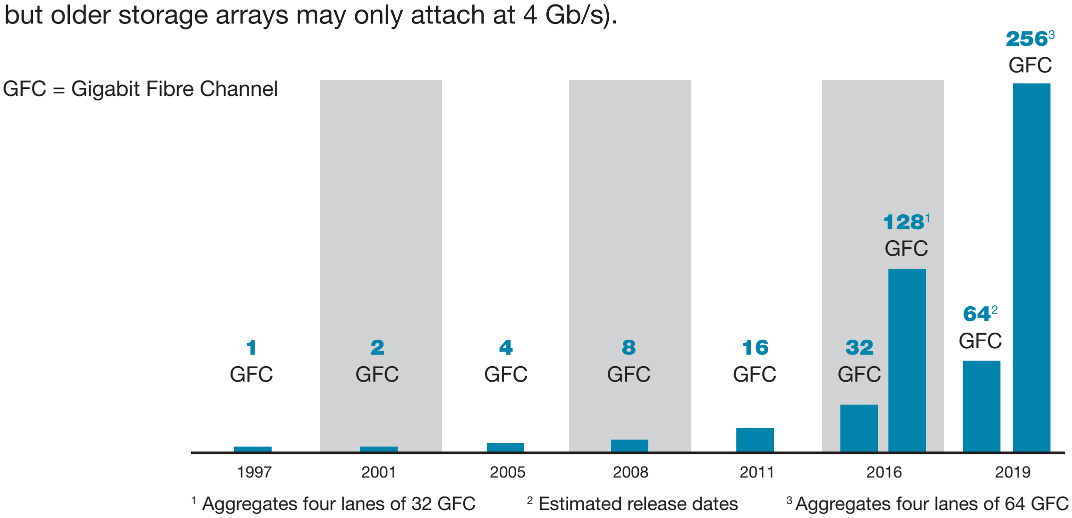

Figure 1. Fibre Channel development since 1997.

PANDUIT®

Data Rate Evolution

Fibre Channel supports the interconnection between servers and storage over high-speed optical fiber cabling. Channel speed evolves in base 2 progression such that there is a doubling of previous data rate generations. Figure 2 summarizes the development of data rates compared to historical and future Fibre Channel implementations and includes Fibre Channel speeds and Ethernet speeds used (and proposed) for FCoE.

THE PAST, PRESENT, AND FUTURE OF FIBRE CHANNEL

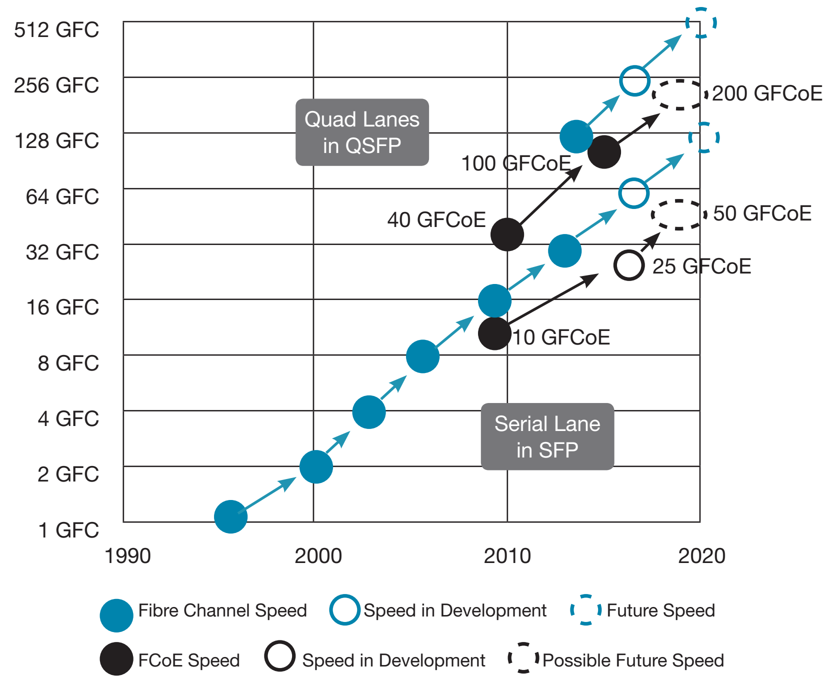

Figure 2. The Fibre Channel Industry Association has developed a roadmap for Fibre Channel data rates in graphical format.

Roadmap for Future Fibre Channel Variants

Fibre Channel ports support a variety of transceivers that have different data rate capability and reach. The Fibre Channel standard historically requires that current transceivers support two backward generations of speeds (e.g., 8 GFC transceivers support 4 GFC and 2 GFC).

Currently, two types of transceiver form factors dominate the market for Fibre Channel switch-to-switch links and switch-to-host/host bus adapter (HBA) links: the small form factor pluggable (SFP, duplex LC) and the Quad SFP (QSFP, 8/12 fiber MPO). In applications where interconnected switches are collocated or hosts/servers are near an edge SAN switch, copper-based direct attach cables (DACs) supporting Fibre Channel may be deployed (typically at less than 10 meters [m]).

SFP-style modules, which are currently shipping in very high volumes, constitute the majority of transceivers for applications from 1-32 GFC. The original design for the SFP was meant to support 5 Gb/s or below, so new variants were designed for applications beyond 4 GFC, such as:

a) SFP+ for 8 GFC, 10 GFC, 10 Gigabit Ethernet (10 GbE) and 16 GFC

b) SFP28 for 32 GFC (and 25 GbE)

Preliminary data suggests that SFP28 modules will support data rates beyond these (namely 50 and 64 GFC), but this is yet to be proven. All generations of these SFP modules have been specified with progressively better signal integrity to match the increasing data rates.

QSFP modules were first deployed in InfiniBand and higher- speed Ethernet applications like 40 GbE. These modules have four optical/electrical parallel lanes supporting higher speeds 100 GbE (4 x 25 Gb) and 128 GFC (4 x 32 G).

QSFP was initially defined to support a line rate of SFP (4 x 5 Gb/s). QSFP style modules constitute many transceivers for 40 GbE through 128 GFC applications. The original QSFP design was meant to support 20Gb/s or below, so new variants were designed for Fibre Channel applications beyond 16 GFC. Much like the SFP series of serial duplex modules for 32 Gb and lower data rate Fibre Channel, new generations of QSFP, QSFP+ and QSFP28 have been developed to satisfy enhanced data rate capabilities:

a) QSFP+ for 128 GFC (with possible breakout functionality for 4 x 32 GFC)

b) QSFP28 for future Ethernet/Fibre Channel (200 GbE and 256 GFC are targets for this)

Future Fibre Channel Technology

To reduce costs where possible, Fibre Channel SAN links utilize shortwave transceivers coupled with multimode fiber (MMF) within the data center. Singlemode fiber (SMF) and transceivers (typically a more expensive option) are deployed for interswitch links (ISL) between data center halls/buildings.

Figure 3 shows the reach of current and future speeds of Fibre Channel over various generations of Optical Multimode (OM) fiber. The second generation of laser-optimized MMF (OM4) is required for 32 GFC links at 100m.

For longer reaches, long-wave transceivers using SMF is required. SMF transceiver variants under the Fibre Channel Physical Interface specification support up to 10 kilometers (km) of fiber (some lower specification/cost variants support 2km).

Like Ethernet 100 Gb parallel singlemode 4-channel (PSM4) low- cost 500m transceivers, 128 GFC will be supported with the same technology as 128 GFC-PSM4 transceivers. Similarly, the technology deployed in 2km 100 Gb CWDM4 for Ethernet will be used for 128 GFC. It is expected that switch manufacturers will use PSM4 technology to implement “break-out” functionality into the PSM4 ports to provision ISL connection between 128 G QSFP28 and 32 G SFP28 modules.

MULTIMODE FIBER LINK DISTANCE (METERS)

Product Naming	Throughput (Mbytes/s)	Line Rate (Gbaud)	T11 Specification Technically Complete (Year)*	Market Availability (Year)*
1 GFC	200	1.0625	1996	1997
2 GFC	400	2.125	2000	2001
4 GFC	800	4.25	2003	2005
8 GFC	1,600	8.5	2006	2008
16 GFC	3,200	14.025	2009	2011
32 GFC	6,400	28.05	2013	2016
128 GFC	25,600	4x28.05	2014	2016
64 GFC	12,800	56.1	2017	2019
256 GFC	51,200	4x56.1	2017	2019
128 GFC	25,600	TBD	2020	Market Demand
256 GFC	51,200	TBD	2023	Market Demand
512 GFC	102,400	TBD	2026	Market Demand
1 TFC	204,800	TBD	2029	Market Demand

GEN 6 Fibre Channel

GEN 7 Fibre Channel

Figure 3. Reach of multimode optical fiber links.

Fibre Channel Cable Infrastructure

IEC and TIA standards define links as the permanent optical fiber cabling infrastructure over which the active equipment must communicate. Links do not include equipment patch cords connecting active network devices in equipment distribution areas or the cross-connect cords.

ISO/IEC and TIA standards define link testing as performance verification of fixed (permanent) segments of installed cabling. Completion of testing ensures that links that pass standards-based (or application-based) limits and can be reliably configured into passing channels by adding good-quality patch cords.

The ANSI T11 Fibre Channel standard defines the channel requirements that form SANs. Each Fibre Channel Physical Media Dependent (PMD) layer has different characteristics, depending on the fiber cable plant deployed. A designer approaching the problem of designing a flexible and robust cabling system is presented with a set of cabling system “dials” that can be manipulated to ensure proper function of the SAN channels.

The application link power budget for Ethernet and Fibre Channel does not include the connectors that are attached to equipment on either end of the link. These are built into the link power budget as the difference between the minimum transmitter power into the fiber and receiver minimum sensitivity. As a result, the number of connectors in the channel is the total number of mated pairs. Connectors that are mated to the optical transceivers are not considered mated pairs.

Application Standards Power Budgets

The overall power budget for an optical channel is determined by the application standard, such as Ethernet, and is based on the magnitude of power penalties, or impairments, as well as the maximum length of the channel. Typically, most of these optical impairments are small, under 0.3 decibels (dB). However, inter-symbol interference (ISI), which is a function of the optical fiber's bandwidth, and connector insertion loss (IL) contribute large optical penalties and are the two primary impairments that limit the reach of the channel. The impact of these impairments is strongly influenced by the quality and practices used in the construction and testing of the channel.

The Fibre Channel standard historically requires that current transceivers support two backward generations of speeds

(e.g., 8 G Fibre Channel transceivers support 4 G Fibre Channel and 2 G Fibre Channel)

There are two sources of loss: IL (loss in the mated connectors), and attenuation (loss of the laser energy within the optical fiber). IL is a critical parameter that determines the performance of a channel. IL in a link depends on the number of optical fiber connections and the individual losses of mated fiber connector pairs therein. The total loss present in the link is the sum of connector/splice losses and the attenuation generated by the fiber length.

In principle, one can trade fiber attenuation for connector IL or ISI power penalties for IL; however, this must be done with caution. Engineered links are those channels designed by making tradeoffs between these parameters.

As an example, consider an OM4 (M5F in Fibre Channel), 16 GFC link with an installed reach of 50m; this is a third of the maximum specified reach of an 150m engineered link (Table 1).

ISI for this channel is significantly less than it is at 150m. As a result, a larger connector IL of 2.4dB can be tolerated. Alternatively, the ISI penalty can be reduced by using the increased fiber bandwidth of OM4.

Table 1. Gb/s Fibre Channel reach/power budget vs. total connector insertion loss

DISTANCE (m)/LOSS BUDGET (dB)

Fiber Type	Connection Loss (dB)			
	2.4	2.0	1.5	1.0
OM2	N/A	25/2.09	35/1.63	40/1.14
OM3	40/2.54	75/2.27	100/1.86	120/1.43
OM4	50/2.58	100/2.36	125/1.95	150/1.54
Signature Core™ OM4	147/2.91	175/2.61	200/2.2	215/1.75

Structured Cabling Model for Fibre Channel

Structured cabling allows for low-risk moves, adds, and changes (MACs) and for efficient jumper management between ISLs, servers, and host storage. This enables dynamic change management in data centers and provides a path for growth and migration to newer technologies. This cabling model enables the highest-density switches and patching to be deployed without sacrificing manageability. Fibre Channel power budgets are trending downwards with the increase in data rates presented by each new generation of Fibre Channel. It is also important that an implementation path of "any-to-any" connectivity is taken for high-density environments. This presents a challenge to the power budget as more mated connector pairs are present in such channels.

It is very important to engage early in the design phase with structured cabling subject matter experts to understand the limitations of the Fibre Channel link model when such engineered links are being considered.

The goal of a structured cabling model is for all ports in an "any-to-any" cross-connect map to be the logical representation of switch, server, and storage ports distributed around the data center. This cross-connect switchboard is described as the main distribution area (MDA) in TIA 942 and as the central patching location in IBM Fiber Transport Systems (FTS). Cross-connects between end points are done by jumper cables between patch panels/enclosures in the MDA. This approach ensures that active equipment is not disturbed except in the case of outage or technology change.

Structured cabling provides a standards-based system that makes documentation of MACs (using TIA 606) manageable and scalable to support future Fibre Channel technology.

Figure 4 on page 9 shows how multi-fiber push on (MPO) trunk cables are connected to MPO/LC cassette modules. MPO terminated trunk cables are built on increments of twelve fibers and can traverse long distances between patch panels/enclosures in the data center. Such trunk cables can be connected to the modules indicated above or connected to other MPO extension trunks via MPO coupler panels. An option for a static environment that does not require patching is to deploy an LC break-out harness to connect to the MPO trunk through an MPO coupler panel.

High-density MPO/LC cassette modules in patch panels/enclosures have up to 72 LC or MPO ports per rack unit (RU). Figure 5 on page 11 provides a description of the elements of plug-and-play optical fiber that support easy implementation of the structured cabling model.

MPO-LC harnesses (breakout cables) can connect SFP+ transceivers on a high-density SAN director (LC end) to an MPO-based fiber trunk.

Modular cassettes and MPO fiber adapter panels support the connector trunk cables terminated with MPO connectors. The trunk cables plug in the back of these components and LC or MPO patch cords connect to the front. Modular plug-and-play architecture minimizes the total installed cost and provides ultimate scalability.

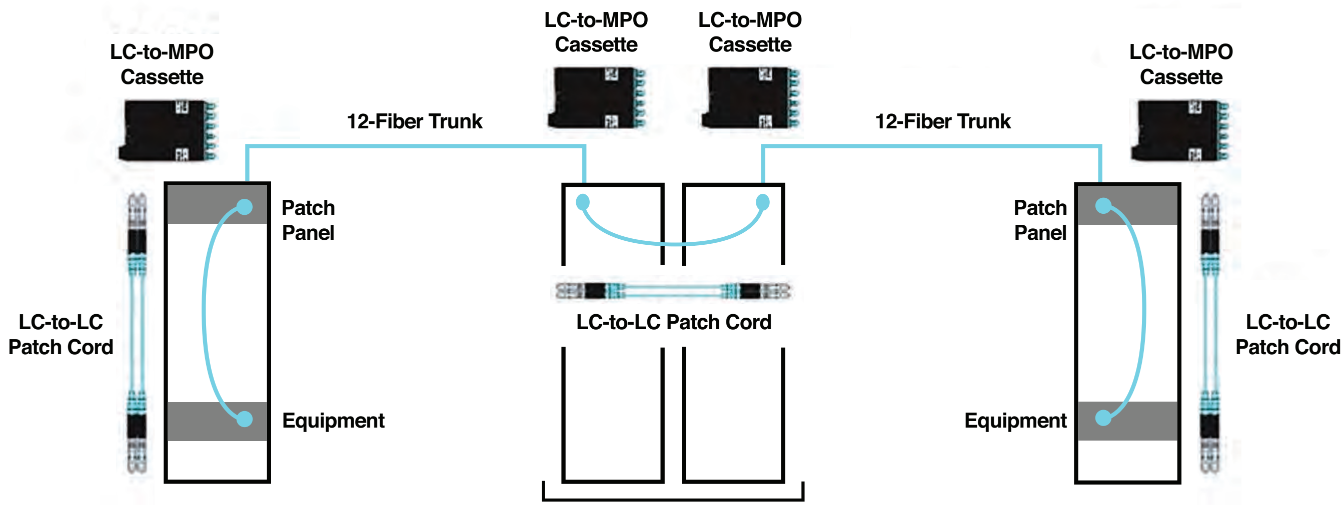

Figure 4. Cross-connect plant model.

Modular cassette-based fiber distribution systems support 72 LC ports (two fiber duplex) in SAN. A SAN rack/cabinet with this system can support 3,024 LC ports. The concern with such high-density solutions in an area such as MDA cross-connect is manageability and risk of disturbing adjacent (potentially active) circuits. Users should look for fiber distribution systems that minimize risks in active patching areas.

Typical Data Center SAN Cabling Systems

Channels can be grouped at the top level by the number of mated connector pairs in the channel. Connections are typically made at a patch panel or enclosure that may serve different purposes (e.g., consolidation, port-mapping, cross-connect) and such panels are either deemed static (as in a zone cable consolidation point under the floor with few MACs) or dynamic (as in a cross-connect that provisions “any-to-any” MACs). Most frequently, Fibre Channel links are deployed in one of three ways:

1. Point-to-point: direct connection between Fibre Channel ports with fiber cabling and no mated pairs of connectors in the channel.
2. Interconnect cabling: A single permanent cabling segment terminated into patch panels or enclosures on each end and jumpers connecting to equipment transceivers on each end.
3. Cross-connect cabling: Two segments of permanent cabling (each with patch panels or enclosures on their ends), cross-connected in a central location and jumpers connecting to equipment transceivers on each end (Figure 5).

Structured cabling uses fiber distribution systems (high-density patch fields) that connect permanent infrastructure (MPO trunks), connecting various equipment distribution areas within the data center. The deployment of pre-terminated MPO systems (rather than field-installed systems) in a cross-connect structured cabling model quickly provisions infrastructure where device ports can be connected to any other port. As networking equipment becomes dense and SAN port counts climb into the thousands, managing connected cables becomes a challenge. Historically, connecting cables

directly to individual ports on equipment was considered manageable for small deployments (point-to-point or interconnect models). Applying the same models to high port-count equipment makes it nearly impossible to perform low-risk MACs.

Many cabling configurations can be supported by Fibre Channel, but implicit knowledge of the power budget for the transceivers being deployed is required, especially in channels that depart from the basic channel shown in the Fibre Channel-PI standard (two mated pairs of fiber connectors at 0.75dB per mated pair).

Large data centers use four cassette module channels having MPO trunk cables that connect disparate equipment distribution area (EDA) patch facilities to the MDA.

Figure 4 shows a four-cassette module channel that uses trunk cables deployed at either side of the cross connect/MDA. Each trunk cable end connects to modular fiber cassettes, presenting a channel that includes four such cassettes as milestones for added connector IL. The LC-to-LC patch cord shown in the MDA connects mapped ports within adjacent trunk cables to each other. LC-to-LC patch cords on either end of the link connect to the transceiver present in the electronics. This scenario consists of two trunk cables, two EDA LC-to-LC patch cords, one MDA LC-to-LC patch cord and four LC cassette modules. An alternate configuration replaces one of the LC cassette modules (typically on the high-density SAN Director end) with an MPO-LC harness and a MPO coupler panel. This configuration also minimizes connector IL as one LC-to-LC mated pair is removed from the channel (no LC patching at the SAN Director EDA).

This four-cassette module channel starts at an SFP+ port on the storage device at the bottom left of Figure 4. An LC-to-LC patch cord connects the storage device to the first cassette that connects to the MPO trunk. This trunk connects to the leftmost cassette in the MDA and an LC-to-LC patch cord connects to the opposing cassette in the MDA and ultimately to the trunk on the right. This trunk connects to the cassette in the switching area. The rightmost LC-to-LC patch cord connects to an SFP+ port on the switch.

This is one example that is a typical deployment of structured cabling infrastructure representing cross-connect MDA.

THE COMPONENTS OF STRUCTURED FIBER CABLING

MPO Trunk Assembly
MPO terminated optical fiber trunk assemblies are typically 12-144 fibers and create the permanent fiber links between patch panels in a structured cabling environment. They are pre-terminated from the manufacturer with MPO connectors at a specified length and have a pulling grip for easy installation.

Fiber Distribution Enclosure
Connector housings are physically mounted in a 19-inch rack or cabinet. They are typically offered in various sizes such as 1 RU, 2 RU, or 4 RU which refers to the amount of rack space required for mounting.

Plug/Play Cassette
MPO to LC cassettes are installed into fiber distribution enclosures. They break out MPOs in the trunk assemblies into LC adapters. Trunk cables plug into the rear of the cassette, and LC jumpers plug into the adapters on the front of the cassette.

QSFP Breakout Cassette
MPO to LC modules are installed into the connector housings. They break out the MPO connection from the trunk cables into LC connectivity. Thus, the trunk cables plug into the rear MPO of the module, and LC jumpers plug into the front of the module.

MPO Fiber Adapter Panel
MPO fiber adapter panels (FAPs) are installed into fiber distribution enclosures. They offer a connection point between the MPO trunks and MPO jumpers or breakout harnesses. Thus, the trunk cables will plug into the rear of the panel, and the MPO jumpers or harnesses will plug into the front of the panel.

MPO-LC Harness (breakout cable)
Harness assemblies are used for breaking out the MPO connector into multiple LC connections. In the case of breaking out a QSFP into four lanes at the QSFP base line rate, these assemblies are wired identically to QSFP breakout cassettes.

MPO or LC Jumpers
MPO jumpers serve to interconnect QSFP active ports, connect cassettes to permanent links, and provide jumper connection between active QSFP ports and structured cabling systems.

Figure 5. The elements of plug-and-play optical fiber that support easy implementation of the structured cabling model.

Fibre Channel Connectivity Choices

While the previous example shows LC-MPO cassettes and MPO trunk cables deployed to form a channel, there are many choices of components in implementing the structured cabling model. Trunk cables terminated with LC connectors and plugged into the back of LC-based adapter panels in fiber distribution enclosures would obviate the need for cassette modules and would serve to reduce the total connector IL in the channel (four mated pairs of MPO connectors are removed). The design of Fibre Channel supports a variety of cable plant components and architectures to meet customer needs.

There are limits to how many connections a link may support. The four-cassette module model shown has higher IL, and such loss needs to be limited to meet the requirements of the Fibre Channel-PI standard. The basic two connector model of the Fibre Channel standard limits the total connector-based IL to 1.5dB (two mated pairs at 0.75dB each), so there are a few approaches to design. The first and simplest is to use lower loss modules so that the sum of connector IL for the four cassettes is less than 1.5dB (0.375dB maximum each). The second demands explicit knowledge of the channel model and involves tradeoffs of power penalties. It is possible to have more connector loss than 1.5dB in the channel by trading power penalties of fiber attenuation and ISI for additional connector loss.

It is very important to engage early in the design phase with subject matter experts at data center structured cabling providers to understand the limitations of the Fibre Channel Link Model when such engineered links are being considered.

Summary

Fibre Channel deploys cost-effective MMF and SMF transceivers to facilitate high-speed SAN ports in data centers that range from enterprise-scale equipment rooms to hyperscale data centers. The majority of Fibre Channel channels deployed use MMF with a cross-connect/central patching location positioned close to large director-class Fibre Channel switches. In larger deployments with high density directors, bringing ports off the switch and mapping them in the cross-connect is best practice.

To achieve the flexibility of “any-to-any” patching in a cross-connect, multiple connector hops will be deployed in the channel. In the four-cassette scenario presented, connector IL must be managed with the knowledge of the power budget available given the desired link reach and fiber type deployed.

For the longest channels within or between buildings, longwave transceivers and SMF fibers must be deployed. In general, there is more power budget available when using SMF solutions (for 10km variants). This allows designers to deploy cabling solutions with much higher overall loss to support these types of channels.

The Fibre Channel roadmap will support customer SAN applications well into the future, moving along by base two increments of data rate. The sixth generation of Fibre Channel has been designed to deliver 128 Gb through four transmit lanes of 32 Gb and four receive lanes of 32 Gb (MPO-based parallel optics cabling). Although not explicit in the Fibre Channel-PI-6P standard, break-out capability to 32 Gb will be supported by OEM vendors of Fibre Channel switch gear. Other lower cost modules based on SWDM (OM5 MMF) and PSM4 and CWDM4 (SMF) will be developed for 128 GFC and Fibre Channel roadmap items such as 256 Gb Fibre Channel (and beyond).

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 WHITE PAPER 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.

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