White Paper

Fiber-Optic Cabling Connectivity Guide for 40-Gbps Bidirectional and Parallel Optical Transceivers

What You Will Learn

As data centers consolidate into more complex systems, they take advantage of new speed increases and technologies. The speed changes from 1- to 10-Gbps infrastructure were easy to understand. With speeds in the data center now increasing from 10 Gbps to 40 Gbps and eventually to 100 Gbps, different optical technologies and cabling infrastructure are required. This document introduces the cable requirements for 40-Gbps infrastructure and fundamental cabling principles supporting Cisco Nexus® 9000 Series Switches. Although alternative cabling options are mentioned (Twinax and active optical assemblies), the main focus of the document is cabling for pluggable optical Enhanced Quad Small Form-Factor Pluggable (QSFP+) modules.

The new Cisco Nexus 9000 Series provides high 1-, 10-, 40-, and (future) 100-Gbps Ethernet densities with outstanding performance and a comprehensive feature set. The Cisco Nexus 9000 Series provides a versatile platform that can be deployed in multiple scenarios - direct-attach 1-, 10-, and 40-Gbps access and collapsed aggregation and access deployments, leaf-and-spine architecture, and compact aggregation solutions.

Structured (Semi-permanent Links) or Unstructured System

When designing a networking system, it is important to plan the cabling system in advance. The goal is to address current network requirements as well as accommodate future growth. A structured cabling system provides a flexible cabling plan to address the commonly performed tasks of moving, adding, and changing the infrastructure as the network grows. A good analogy for a structured cabling system is the electrical wiring in your home. When connecting appliances and devices, you require only a 5-foot connection to the closest electrical outlet. However, without an electrical outlet, all appliances would have to connect directly to the breaker or panel, requiring a cable of 200 feet or more. This approach would be inefficient and would become unmanageable as you add multiple appliances and devices throughout the home. The structured wiring is those permanent wires in the walls and ceilings to electrical outlets throughout your home. The same concept applies to data centers with the deployment of robust cabling links between two patch panels (similar to the outlets in the wiring example). Structured cabling becomes a necessity as the infrastructure grows and as constant moves and changes reinforce the need for a reliable network that is also easy to troubleshoot.

Structured cabling requires additional initial investment to create the cabling infrastructure, but the recurring benefits more than outweigh the slight additional incremental cost. Imagine the cost of deploying a two-fiber optical jumper each time a new server is placed in the data center. Further, regardless of whether the data center has a raised floor or uses overhead cabling, both result in time-consuming and inefficient deployment in an unstructured environment. Likewise, management of such an environment is cumbersome, increasing the risk of outages caused by human errors.

Structured cabling uses fiber termination connector panels that are connected through permanent links of optical cabling, typically configured in a star topology. All cabling in the data center coming from the server area is consolidated in a central location near the core, aggregation-layer, or spine switch in the network (analogous to the breaker or power panel in the home electrical system analogy). The permanent pre-terminated trunk cables branch to the zones in the data center, which contain servers, storage, or network devices. Note that with structured cabling, you still need some device-to-device connections at the access layer. As you can see in Figure 1, when you make these short connections within the same cabinet or even a few cabinets away, patch panels may not be required. Likewise, patch panels would not be required for inter-switch link connections.

Figure 1. Structured Cabling with Patch Panels

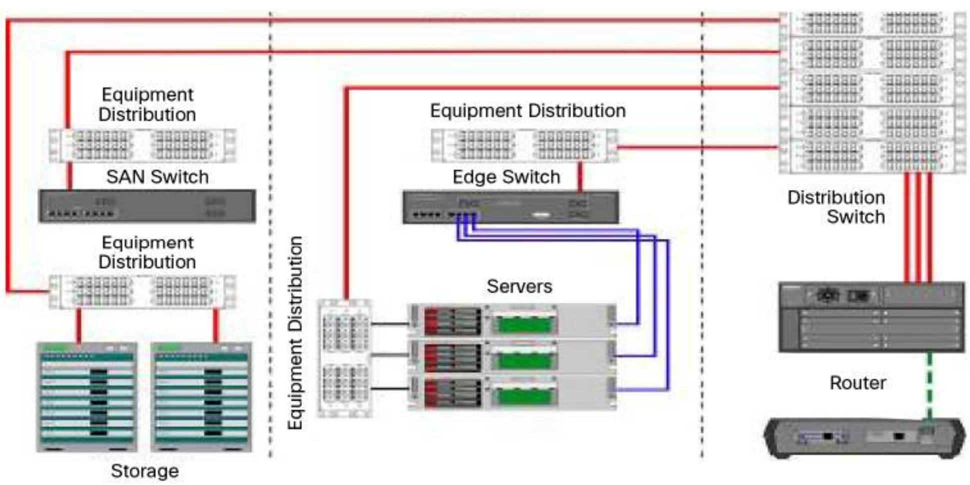

Unstructured cabling occurs when optical links are deployed point to point or device to device with no patch panels installed in the link. In this situation, cabling pathways become congested with an entangled mess of two-fiber optical patch cords (Figure 2). Likewise, routing new patch cords in ceiling or floor trays all the way across a data center each time a new device is deployed is extremely inefficient.

Figure 2. Unstructured Cabling

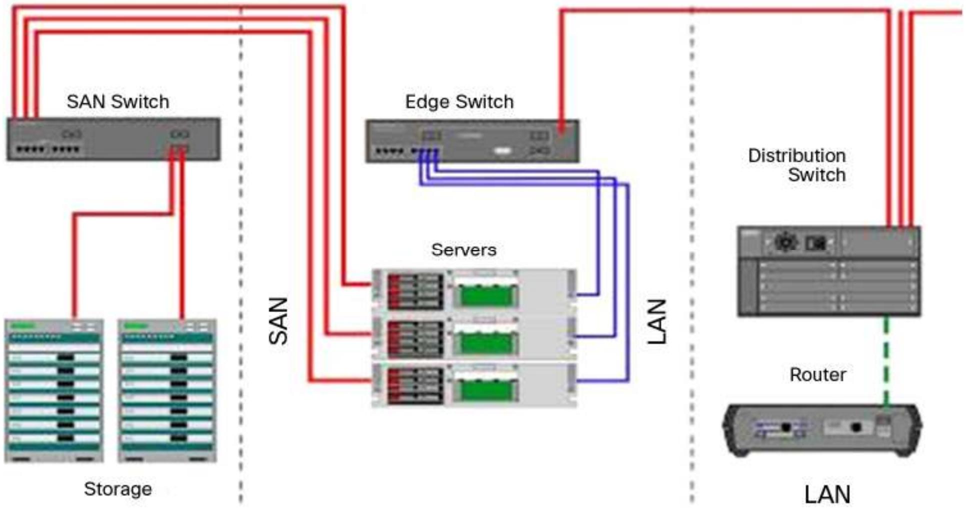

Types of Fiber Optic Fiber

As a result of the emergence of high-data-rate systems such as 10, 40, and 100 Gigabit Ethernet, laser-optimized multimode fiber (MMF) has become the dominant fiber choice. These 50-micron fibers are optimized for the 850-nanometer (nm) transmission of vertical-cavity surface-emitting laser (VCSEL)-based transceivers. The TIA-492AAAC OM3 detailed fiber standard was released in March 2002, and the TIA-492AAAD OM4 detailed fiber standard was released in August 2009. Corning Cable Systems suggests installing either OM3 or OM4 cabling in the data center space depending on length requirements. The two fibers have different bandwidths (information carrying capacity), which results in different achievable lengths for the same transceivers. Table 1 shows the achievable distances based on the OM3 and OM4 fibers at various data rates. TIA 942-A Telecommunication Infrastructure Standards for Data Centers recognizes only the OM3 and OM4 MMF types (it removed OM1, which was 62.5 microns, and OM2, which was the standard 50 microns). The standard also provides guidance to recommended OM4 cables using LC and MTP connectivity.

Table 1. 850-nm Ethernet Distance (m)

	1-Gbps 1000BASE-SX	10-Gbps 10GBASE-SR	40-Gbps 40GBASE-SR-BD	40-Gbps 40GBASE-SR4	40-Gbps 40GBASE-cSR4
OM3 (50 microns)	• 1000m • 2 fibers	• 300m • 2 fibers	• 100m • 2 fibers	• 100m • 8 fibers	• 300m • 8 fibers
OM4 (50 microns)	• 1000m • 2 fibers	• 400m • 2 fibers	• 125m • 2 fibers	• 150m • 8 fibers	• 400m • 8 fibers

Types of Optical Transceiver Modules

The transceiver is an electronic device that receives an electrical signal, converts it into a light signal, and launches the signal into a fiber. It also receives the light signal, from another transceiver, and converts it into an electrical signal. For data rates greater than or equal to 1 Gbps, a multimode transceiver uses an 850-nm VCSEL transceiver.

SFP+ is the dominant transceiver form factor used for 1 and 10 Gigabit Ethernet applications. The transceiver uses an LC optical connector interface. For more information, see http://www.cisco.com/en/US/prod/collateral/modules/ps5455/data_sheet_c78-455693.html.

The QSFP+ transceiver is the dominant transceiver form factor used for 40 Gigabit Ethernet applications. In 2010 the IEEE standard 802.3ba released several 40-Gbps based solutions, including a 40GBASE-SR4 parallel optics solution for MMF. Since then, several engineered solutions have been released, including 40GBASE-CSR4, which is similar to 40GBASE-SR4 but extends the distance capabilities. Another solution released by Cisco is a bidirectional 40-Gbps transceiver that uses a two-fiber LC optical interface. For more information, see http://www.cisco.com/en/US/prod/collateral/modules/ps5455/data_sheet_c78-660083.html.

Table 2 summarizes the differences between QSFP transceiver optical interfaces and electrical interfaces. Table 3 summarizes the connectivity options.

Table 2. Comparison of QSFP Transceiver Optical Interface and Electrical Interface

Description	Electrical and Optical Lanes Diagram	Physical Transceiver
Parallel optical transceivers used for 40GBASE-SR4 and 40GBASE-CSR4 have 10-Gbps electrical lanes that are mirrored in the optical outputs and thus require eight fibers with an MTP connector interface. Each fiber either transmits (Tx) or receives (Rx) 10-Gbps traffic at a single wavelength.	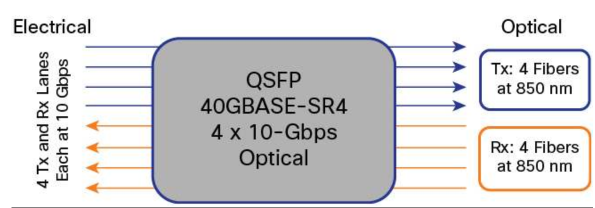	
Bidirectional optical transceivers used for 40GBASE-SR-BD use the same 10-Gbps electrical lanes, which are then combined in the optical outputs, thus requiring two fibers with an LC connector interface. Each fiber simultaneously transmits and receives 20-Gbps traffic at two different wavelengths.	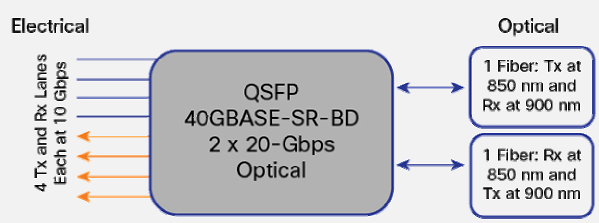	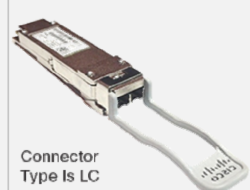

Table 3. Connectivity Options

	Example	Advantages	Disadvantages
Pluggable optical modules	• QSFP-40G-SR-BD • QSFP-40G-SR4 • QSFP-40G-CSR4 • QSFP-40GE-LR4	Allows extended-reach capabilities (up to 400m on MMF and 10 km on single mode fiber [SMF]); cable links and optical engines are separate and thus can be upgraded independently	More expensive than other short-reach direct-attach options
Passive and active direct-attach copper	• QSFP-H40G-CU1M • QSFP-4SFP10G-CU1M • SFP-H10GB-CU2M QSFP-H10G-ACU7M	Allows lower-cost short-reach capability; typically used for top-of-rack (ToR)-to-server connectivity	Limited to 5m (passive) to 10m (active), and coaxial cable can be stiff and bulky in high-density deployments
Active optical cable (AOC) assemblies	QSFP+-AOC (1, 2, 3, 5, 7, and 10m)	Allows lower-cost short-reach capability with more flexible cabling; typically used for ToR-to-server connectivity	Limited to less than 10m; reconfigurations of length or failed transceiver requires replacement of entire assembly

Typical Optical Cabling Components

Table 4 provides an overview of typical cabling components. These components will be mentioned in the cabling discussions throughout the remainder of this document.

Table 4. Typical Cabling Components

Component	Product Photo	Description
MTP trunk cable		These fiber trunk cables are typically 12 to 144 fibers and create the permanent fiber links between patch panels in a structured environment. They are pre-terminated from the manufacturer with MTP connectors at a specified length and have a pulling grip for easy installation.
Connector housings		The connector housings are physically mounted in a 19-inch rack or cabinet. They are typically offered in various sizes such as 1 rack unit (1RU), 2RU, or 4RU, which refers to the amount of rack space required for mounting.
MTP-LC module		The MTP-LC module is installed in the connector housings. The module breaks out the MTP connections from the trunk cables to LC connections for 1- and 10-Gbps applications. The trunk cables plug into the rear MTP of the module, and the LC jumpers plug into the front of the module.
MTP connector panel	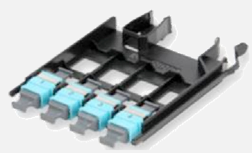	The MTP-MTP connector panel (sometimes called the bulkhead) is installed in the housings. It offers a connection point between the MTP trunk and the MTP jumper for 40- and 100-Gbps applications. The trunk cables plug into the rear of the panel, and the MTP jumpers plug into the front of the panel.
2x3 conversion module		The 2x3 conversion module is used for 4-channel (8-fiber) parallel optics applications, such as 40GBASE-SR4. It allows 100% utilization of trunk cables by converting a pair of 12-fiber MTP connections into three 8-fiber MTP connections. The trunk cables plug into the rear MTP of the module, and the MTP jumpers plug into the front of the module to make the connection to the switch.
LC or MTP jumpers		The LC or MTP jumpers create the connection between the device port and the structured cabling through the connector panel.

Learn more about each of these products in the Corning product catalog; see “Indoor Preterminated Systems” at http://catalog.corning.com/CableSystems/en-US/Default.aspx.

Cisco Nexus 9000 Series Architecture

Today data centers are facing unprecedented requirements for performance, scalability, and agility as virtualization, application workload distribution and mobility, and cloud and big data applications are becoming the main factors influencing data center operation. The Cisco Nexus 9000 Series consists of the Cisco Nexus 9500 platform modular switches and the Cisco Nexus 9300 platform fixed-port switches (Table 5). It provides the next-generation solution to overcome these data center challenges with a high-performance, low cost, flexible and agile application centric infrastructure.

Table 5. Cisco Nexus 9000 Series Switches

Component	Product Photo	Description
Cisco Nexus 9500 platform		The Cisco Nexus 9500 platform consists of modular switches. The Cisco Nexus 9508 Switch is the first switch released for this platform. With more than 30 terabits per second (Tbps) of backplane bandwidth, the switch supports 1, 10, 40, and (future) 100 Gigabit Ethernet interfaces through a comprehensive selection of modular line cards. Configurable with up to 1152 10 Gigabit Ethernet or 288 40 Gigabit Ethernet ports, the switch provides sufficient capacity for both access- and aggregation-layer deployments.
Cisco Nexus 9396PX Switch	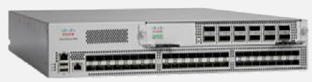	The Cisco Nexus 9300 platform consists of fixed-port switches designed for ToR and middle-of-row (MoR) deployment in data centers. The Cisco Nexus 9396PX is a 2RU non-blocking Layer 2 and 3 switch with 48 1- and 10-Gbps SFP+ ports and 12 40-Gbps QSFP+ ports.
Cisco Nexus 93128TX Switch		The Cisco Nexus 9300 platform consists of fixed-port switches designed for ToR and MoR deployment in data centers. The Cisco Nexus 93128TX is a 3RU 1.28-Tbps Layer 2 and 3 switch with 96 1/10GBASE-T ports and 8 40-Gbps QSFP+ ports.

Cisco Nexus 9000 Series Switches can run in two operating modes: the standard Cisco Nexus device mode with enhanced Cisco NX-OS Software as the operating system, or the Cisco Application Centric Infrastructure (ACI) mode to take full advantage of an automated policy-based approach to system management.

While operating in the standard Cisco Nexus device mode, Cisco Nexus 9000 Series Switches can be deployed in a variety of data center network designs. Cisco Nexus 9500 platform switches, with their high-density line-rate 40-Gbps line cards, can be placed at the aggregation layer to provide aggregated 40-Gbps connectivity for the access switches. Cisco Nexus 9500 platform switches, with their 1- and 10-Gbps line cards, can be deployed as high-performance access-layer end-of-row (EoR) or MoR switches for 1- and 10-Gbps server connectivity. Cisco Nexus 9300 platform switches are well designed for the access layer as ToR switches. The Cisco Nexus 93128TX is also a good choice as a MoR access switch.

As shown in Figure 3, Cisco Nexus 9500 and 9300 platform switches are deployed in the traditional three-tier design. The Cisco Nexus 9500 platform switches are placed at the core aggregation tiers, and the Cisco Nexus 9500 or 9300 platform switches are deployed at the access tier.

Figure 3. Traditional Three-Tier Design

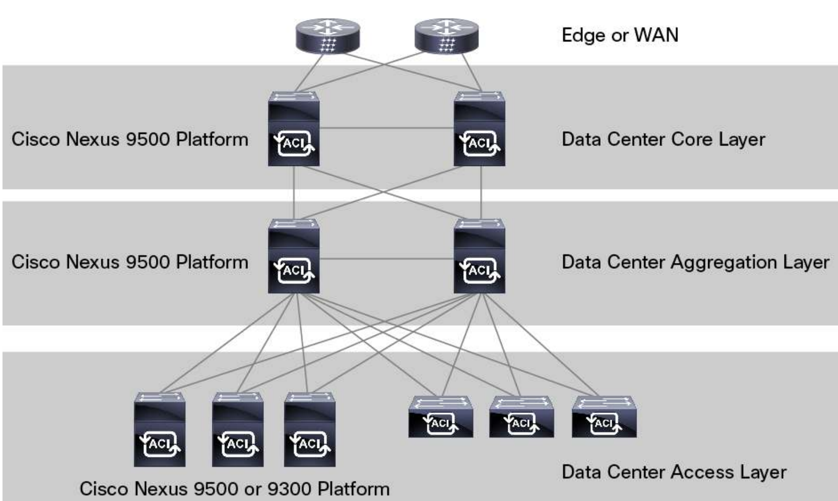

Cisco Nexus 9500 and 9300 platform switches both support Cisco Nexus 2000 Series Fabric Extenders (FEXs). By using the Cisco Nexus 2000 Series, Cisco Nexus 9500 and 9300 platform switches can build a cost-effective and scalable collapsed aggregation and access layer for 1- and 10-Gbps server connectivity with 40-Gbps uplinks to the network. Figures 4 and 5 show Cisco Nexus 9500 and 9300 platform switches, respectively, in this collapsed two-tier design.

Figure 4. Cisco Nexus 9500 Platform Switches and Fabric Extenders for Collapsed Aggregation and Access Layer

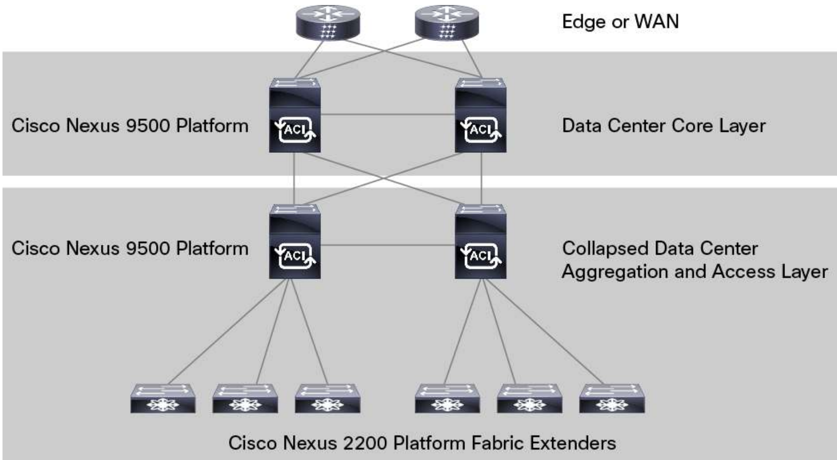

Figure 5. Cisco Nexus 9300 Platform Switches and Fabric Extenders for Collapsed Aggregation and Access Layer

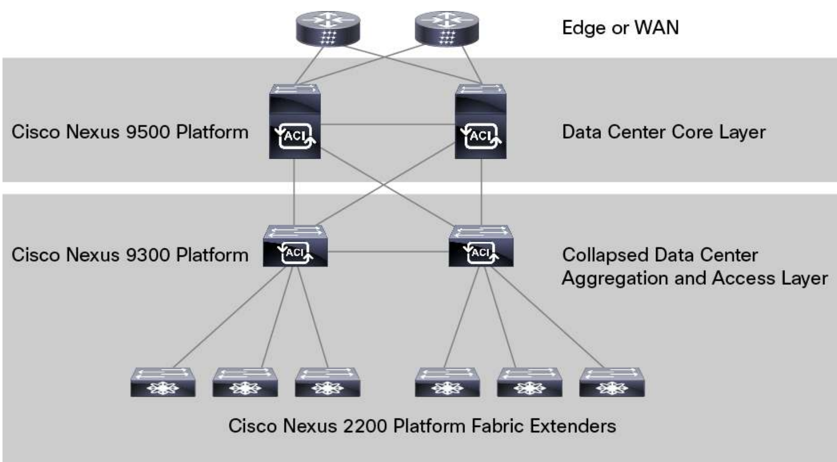

Modern data center applications, such as big data and high-frequency trading (HFT) applications, and virtualized and clustered environments have shifted the data center traffic load from north-south client-server traffic to east-west server-to-server traffic. Data center architects are starting to gravitate to the spine-leaf topology, which flattens the network to two tiers and increases traffic-forwarding efficiency for the increasing amount of east-west traffic. With their non-blocking architecture and high 10- and 40-Gbps port densities, the Cisco Nexus 9500 and 9300 platforms are excellent choices for a spine-and-leaf network design. Figure 6 shows a sample spine-and-leaf network constructed with Cisco Nexus 9500 and 9300 platform switches. With their versatile line-card options, Cisco Nexus 9500 platform switches can be deployed at both the spine and leaf layers.

Figure 6. Cisco Nexus 9500 and 9300 Platforms for Spine-and-Leaf Designs

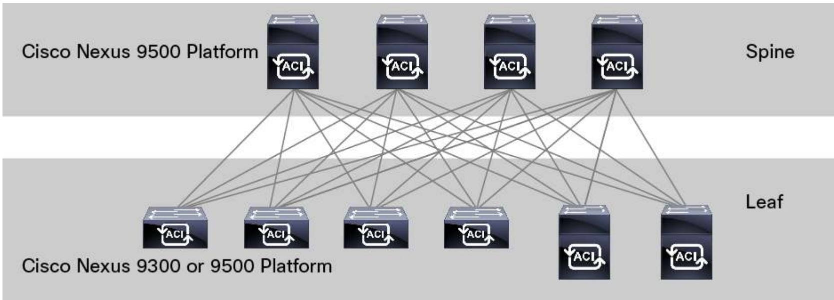

All the preceding deployment scenarios use 1-, 10-, and 40-Gbps cabling for physical connectivity. 1- and 10-Gbps cabling infrastructure is well understood and commonly deployed, but 40-Gbps cabling remains a new challenge to data center operators because of its different transceiver technologies and cabling requirements. The rest of this document provides guidance and options for 40-Gbps fiber cabling designs.

Cabling Options for 40-Gbps Two-Fiber Bidirectional Optics (QSFP-40G-SR-BD)

Since the introduction of 40 Gigabit Ethernet with the IEEE 802.3ba standard in 2010, the only multimode solution on the market has used parallel optics. The parallel optics solution, 40GBASE-SR4, uses eight fibers to transmit four duplex channels each using 10 Gigabit Ethernet. This solution allows an economical path to 40 Gigabit Ethernet data rates, while using many of components of 10 Gigabit Ethernet market solutions.

The change from two-fiber serial transmission to parallel transmission required some changes in the cabling. First, the connector type was converted from the traditional 2-fiber LC duplex connector to a 12-fiber MTP connector. This change created some new challenges: in particular, with the pinning and polarity of the connector. The traditional LC duplex connector uses a ceramic ferrule on each connector, which is aligned in an adapter panel with the use of a ceramic alignment sleeve. However, the MTP connector uses a pinned and non-pinned connector alignment system, making it imperative to always maintain the correct pinning. Likewise, polarity correction of a 2-fiber system can be easily achieved by flipping the position of the LC connector in the duplex clip. However, correction of polarity in a 12-fiber MTP connector can be more challenging because all 12 fibers are in a single ferrule.

In response to these challenges, Cisco developed a two-fiber 40-Gbps Bidirectional (BiDi) multimode solution. This solution uses two different transmission windows (850 and 900 nm) that are transmitted bidirectionally over the same fiber. This approach allows the use of the same cabling infrastructure for 40 Gigabit Ethernet as was used for 1 and 10 Gigabit Ethernet. The pluggable bidirectional transceiver has the same QSFP+ format as the existing 40GBASE-SR4 transceivers. Therefore, the same switch line card with QSFP+ ports can support either parallel optics 40GBASE-SR4 or bidirectional optics 40GBASE-SR-BD solutions.

Thus, when directly connecting a 40 Gigabit Ethernet bidirectional transceiver to another bidirectional transceiver, a Type A-to-B standard LC duplex patch cord can be used. As defined in TIA-568-C.3, a Type A-to-B duplex patch cord is constructed with one (blue) fiber in connector position A on one end and in connector position B on the other end, and likewise for the second (orange) fiber, as shown in Figure 7. This reverse fiber positioning allows a signal to be directed from the transmit position on one end of the network to the receive position on the other end of the network.

Figure 7. Type A-to-B Jumper Fiber Mapping (per TIA-568-C.3)

This type of direct connectivity is suggested only within a given row of cabinets. The jumper assembly is tested only to the requirements of an interconnect cable, as defined in ANSI/ICEA S-83-596-2001. It has less robustness (less tensile strength, less crush and impact resistance, etc.) than a distribution-style trunk cable. Figure 8 depicts a situation in which two 40 Gigabit Ethernet bidirectional ports on two switches are directly linked using a Type A-to-B LC duplex jumper.

Figure 8. Direct Connection for 40 Gigabit Ethernet Bidirectional Transceiver

However, when considering structured cabling, you must consider deployment of more permanent links. The simplest structured cabling link includes a patch panel on both ends of the link, with a jumper assembly making the connection to the electronic ports. This type of cabling is called an interconnect, because both jumpers in the link connect from the structured cabling patch panel to the electronics ports. Figure 9 shows an interconnect link between two bidirectional ports installed in a switch. The link consists of an MTP-based trunk, MTP-LC modules, and LC jumpers. By installing the MMF MTP assembly, you can provide more scalability to accommodate future data rates that may require parallel optics for transmission. This future migration can be accomplished simply by changing the patch panels on each end of the link, without the need to disrupt the cabling infrastructure.

Figure 9. Interconnect for 40 Gigabit Ethernet Bidirectional Transceiver

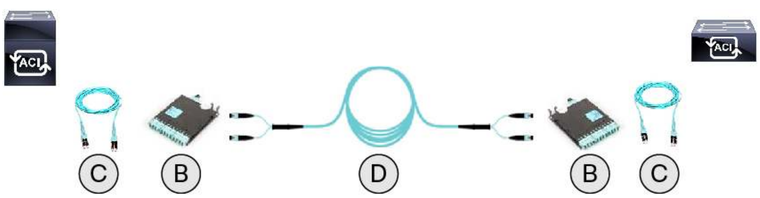

The final cabling approach to consider is a cross-connect design, as shown in Figure 10. In this scenario, two separate structured cabling links connect the two switches through a centralized cross-connect. The advantage of this approach is that it allows the most flexible network configuration. The electronics can be placed in various locations throughout the data center, with structured cabling links between the cross-connect location and designated zone cabinets. When new equipment is installed, only patch cords are required to make the connection from the equipment to the patch panels. Moreover, any port-to-any port connectivity can be achieved at the cross-connect location. This connection is achieved in an orderly and manageable way, unlike what can occur with a direct connectivity scheme in which patch cord assemblies are used to make all port-to-port connections in the data center without structured cabling.

Figure 10. Cross-Connect for 40 Gigabit Ethernet Bidirectional Transceiver

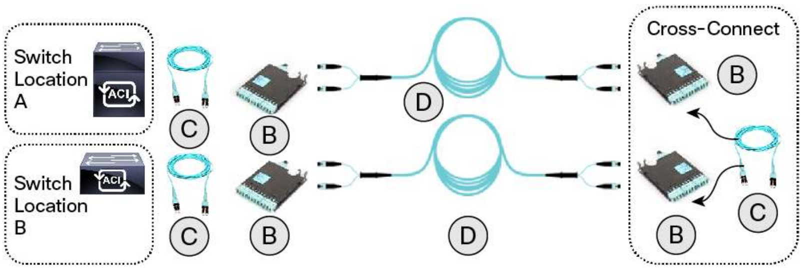

Table 6 provides samples of the part numbers needed for the bidirectional cabling links.

Table 6. Sample of Part Numbers for Bidirectional Cabling Links

Identifier	Part Number	Description
B	ECM-UM12-05-93T	12-fiber LC duplex to MTP (pinned); OM3
C	797902TD120002M	LC uniboot OM3 jumper; 2m
D	G757524TPNDDU100F	24-fiber MTP OM3 trunk cable; 100 ft **Note:** This example uses a 24-fiber trunk cable, but trunk cables can have 12 to 144 fibers at any length.

Cabling Options for 40-Gbps Parallel Optical Modules

As previously mentioned, in 2010 IEEE 802.3ba approved the 40GBASE-SR4 physical-medium-dependent (PMD) multimode parallel optic solution, which uses eight fibers to transmit four duplex channels each at 10 Gigabit Ethernet. As shown in Figure 11, each transceiver transmits over four fibers and also receives transmissions over four fibers.

Also as previously mentioned, parallel optics does require a change from traditional cabling methods, which requires learning and so creates an incentive to move to the bidirectional solution at 40 Gigabit Ethernet. The main advantage of the parallel optics transceiver over the bidirectional transceiver at 40 Gigabit Ethernet is reach. For example, if you cable your data center with OM3 fiber at 10 Gigabit Ethernet, you can support distances up to 300m. Then if you move to 40 Gigabit Ethernet, you can support the same 300m distance with the same OM3 fiber and a 40GBASE-CSR4 transceiver. However, if your cabling distances do not justify the extra distance capability, then the bidirectional solution would be used (Figure 11).

Figure 11. 40GBASE-SR4 and 40GBASE-cSR4 Lane Assignments

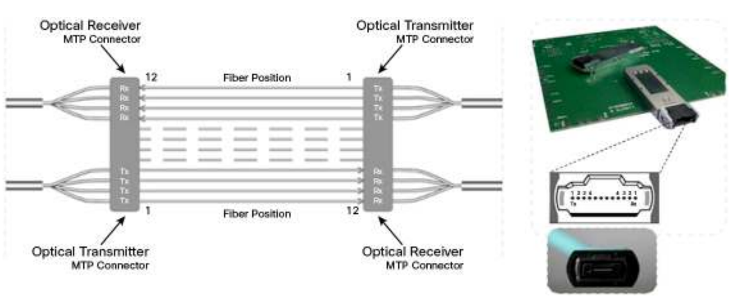

The dilemma is that MTP cable assemblies, which have been used for more than a decade for cabling in the data center, are built on 12-fiber position connectors. Thus, each link has four unused fibers. There are several basic cabling options for parallel optics connectivity. One approach is to ignore the unused fibers and continue to deploy 12 fibers. Another approach is to use a conversion device to convert two 12-fiber links into three 8-fiber links. Three solutions exist (summarized in Table 7 and Figure 12).

- Solution 1: The no-conversion scenario retains the whole 12-fiber based cabling system, but 33 percent of the fiber is not used. Additional cost is associated with the purchase of additional fiber, and your system includes unused fiber.
- Solution 2: This solution uses conversion modules to convert the unused fibers into useable optical links. For every two 12-fiber MTP connectors in the backbone cable, you can create three 8-fiber links. There is an additional cost for the additional MTP connectivity, but that is offset by the cost savings from 100 percent fiber utilization in the structured cabling. When you reuse existing deployed MTP cabling, you gain great value when you use the conversion module to use all previously deployed fiber, and you eliminate the cost of having to deploy additional cabling.
- Solution 3: This scenario uses standard MTP patch panels with a conversion assembly (two 12-fiber MTP connectors on one end going to three 8-fiber MTP connectors on the other end). This approach does not add any connectivity to the link, and it achieves full fiber utilization. However, although this approach may appear attractive, it involves considerable cabling challenges. For example, if you need only two 40-Gbps connections to a piece of equipment, what do you do with that third 8-fiber MTP connection? What if the 40-Gbps ports are in different chassis blades or completely different chassis switches? The result will be long assemblies, which will be difficult to manage in an organized way. For this reason, Solution 3 is expected to be the least desirable and so the least deployed method.

Table 7. Three Cabling Solutions for 40-Gbps Connectivity

	Description	Advantage	Disadvantage
Solution 1: **No conversion**	Uses traditional 12-fiber MTP connectivity and ignores unused fiber	Simplicity and lowest link attenuation	Does not use 33% of the installed fiber, and thus requires more cable raceway congestion
Solution 2: **Conversion module**	Converts two 12-fiber links to three 8-fiber links through a conversion patch panel	Uses all backbone fiber and creates a clean, manageable patch panel with off-the-shelf components	Entails additional connectivity costs and attenuation associated with the conversion device
Solution 3: **Conversion assembly**	Converts two 12-fiber links to three 8-fiber links through a conversion assembly and standard MTP patch panels	Uses all backbone fiber with additional connectivity	Creates cabling challenges with dangling connectors and non-optimized-length patch cords that require customization

Figure 12. Cabling Solutions for 40-Gbps Connectivity

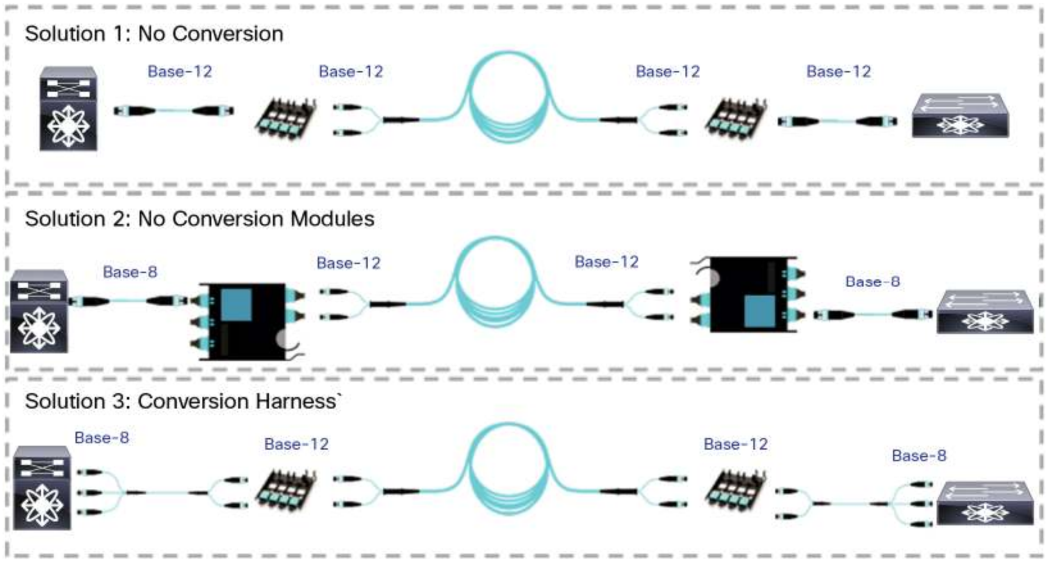

Corning offers components to build all three solutions. However, Corning suggests implementation of the conversion module solution, especially if you are using previously installed MTP trunks. This solution allows 100 percent fiber utilization while maintaining any port-to-any port patching. If you are installing all new cabling, then you could consider the no-conversion solution, assuming that the cable raceway is not a concern. Typically, the conversion harness is deployed only in specific applications, such as at the ToR switch, where 40-Gbps ports are in a close cluster and patching between blades in a chassis switch is not required.

When directly connecting a parallel optics 40 Gigabit Ethernet transceiver to another 40 Gigabit Ethernet transceiver, a Type-B pinless-pinless MTP jumper should be used. As shown in Figure 13, a Type-B MTP jumper assembly, as defined in TIA-568-C.3, has the blue fiber 1 assembled in connector position 1 on one end of the assembly, and this same fiber assembled in connector position 12 on the other end of the assembly. This reverse fiber positioning allows the signal to flow from transmission on one end of the link to reception on the other end.

Figure 13. Type-B Array Patch Cord (per TIA-568-C.3)

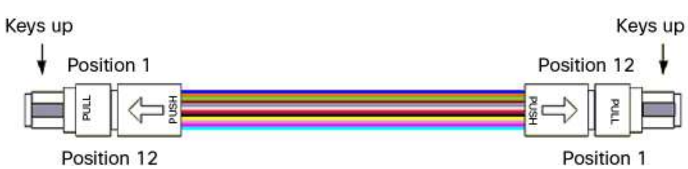

This type of direct connectivity is suggested only within a given row of cabinets. The jumper assembly is tested only to the requirements of an interconnect cable, as defined in ANSI/ICEA S-83-596-2001. It has less robustness (less tensile strength, less crush and impact resistance, etc.) than a distribution-style cable, which would be used for structured cabling trunks. Figure 14 shows two switch ports directly cabled with an MTP jumper patch cord.

Figure 14. Direct Connection for 40 Gigabit Ethernet Parallel Optic Transceiver

Similar to the bidirectional cabling approach, the most basic structured cabling solution is an interconnect. The only difference between an interconnect solution and parallel optics is that the connector type of the patch panels instead is MTP. Figure 15 shows several interconnect link scenarios with various patch-panel options. As previously discussed, the 2x3 conversion modules, depicted in Figure 15a, allow 100 percent fiber utilization and constitute the most commonly deployed method. Another advantage of the conversion module is reduced jumper complexity. Notice that a G jumper, which has a Type-B polarity and is pinless, is used to directly connect two parallel optics transceivers. That same jumper is used on both ends of the interconnect link, thus eliminating concerns about correct pinning.

In Figure 15b, the same trunk is used, but the jumper type is now labeled F. In the bill-of-materials (BoM) shown in Table 8, you can see that an F jumper still has a Type-B polarity, but on one end the MTP is pinned and on the other end a pinless MTP is used. Thus, when you install the jumper, you would install the pinned end in the patch panel, and you would install the pinless end in the electronics. However, because of the pinning, this same jumper could not be used to make a direct connection between two ports. Hence, you can see the advantages of the conversion modules, which both use all the fiber and allow a single-jumper solution.

The combined solution shown in Figure 15c might be deployed when cabling between a spine switch, where the module is placed, and a ToR leaf switch, where the conversion harness and panel are located. The QSFP ports on the leaf switch are closely clustered, so the short breakouts of the 2x3 harness assembly should not be a concern. However, use of the 2x3 harness assembly at the core spine switch is not desirable because patching across blades and chassis is a common practice.

Figure 15. 40 Gigabit Ethernet Parallel Optics Interconnect Link with (a) Conversion Modules, (b) No Conversion, and (c) Combined Conversion Module and Harness Solution

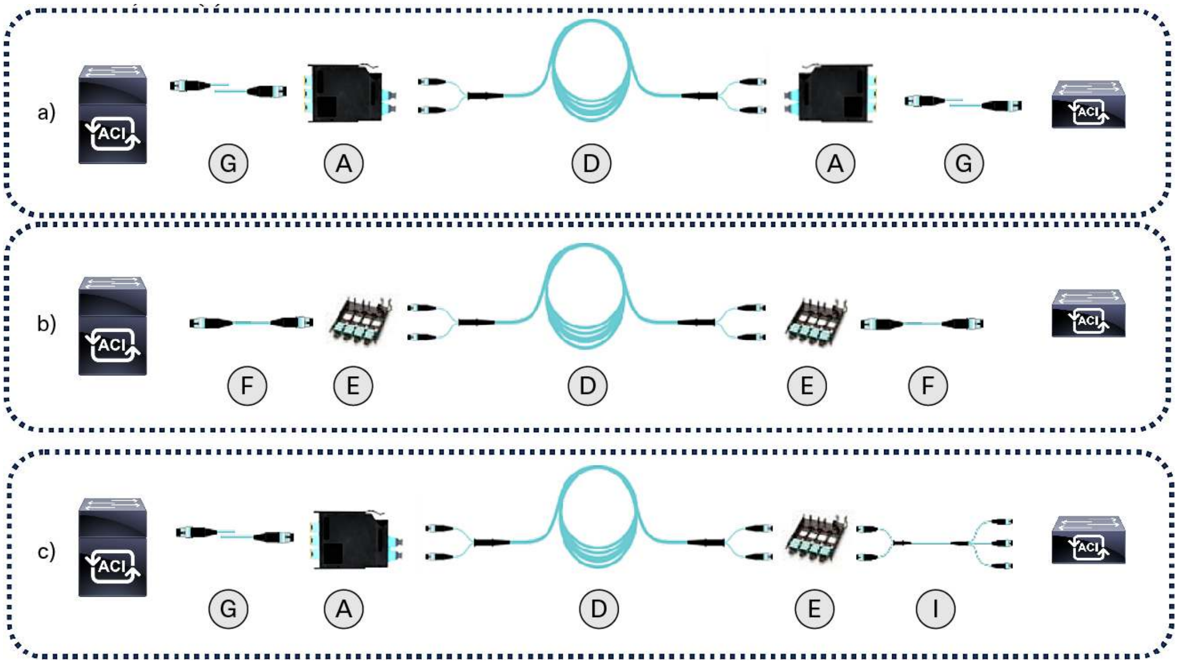

Table 8. Sample Part Numbers for Parallel Optics Cabling Options

Identifier	Part Number	Description
A	ECM-UM24-93-93Q	2x3 conversion module; MTP (pinned) to MTP (pinned) and OM3/4 cable **Note:** A higher-density 4x6 module is also available.
D	G757524TPNDDU100F	24-fiber MTP OM3 trunk cable; 100 ft **Note:** This example uses a 24-fiber trunk cable, but trunk cables can have 12 to 144 fiber at any length.
E	EDGE-CP48-E3	48-fiber (4-port) MTP adapter panel
F	J937512TE8-NB010F	MTP (pinned) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft
G	J757512TE8-NB010F	MTP (pinless) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft
I	H937524QPHKLZ010F ~ H937524QPHKLZ300F	2x3 conversion harness assembly; 12-fiber MTP connections are pinned, and 8-fiber MTP connections are pinless; 24-inch breakout legs; OM3/4 cable; 10~300 ft

As with bidirectional cabling, a cross-connect design allows the most network flexibility. Figure 16 shows two cross-connect network link designs for cabling a 40 Gigabit Ethernet parallel optics transceiver. Figure 16a shows a conversion module example, which again is the most common and preferred method. Notice in this design that all three jumpers (two at the electronics on the left side of the figure and the one at the cross-connect on the right side of the figure) in the link are G jumpers, which according to the BoM in Table 9 are Type-B polarity, and both MTP cables are pinless. Thus, in a conversion module deployment, only one jumper type is used for a direct-connect, interconnect, or cross-connect cabling scenario.

However, notice in Figure 16b that this is not the case for a non-conversion cabling scenario, in which standard MTP patch panels are deployed. Here the patch cords at the electronics are pinless (into the electronics) to pinned (into the patch panel), although the patch cords at the cross-connect are both pinned going into the patch panel. Thus, for a direct-connect, interconnect, and cross-connect cabling scenario, three different pinned jumpers are required.

An alternative approach is to install pinned MTP trunks in the structured cabling, but this approach can be used mainly in new installations because the traditional MTP trunks installed over the past decade have been pinless.

Figure 16. 40 Gigabit Ethernet Parallel Optics Cross-Connect Link with (a) Conversion Modules and (b) No Conversion

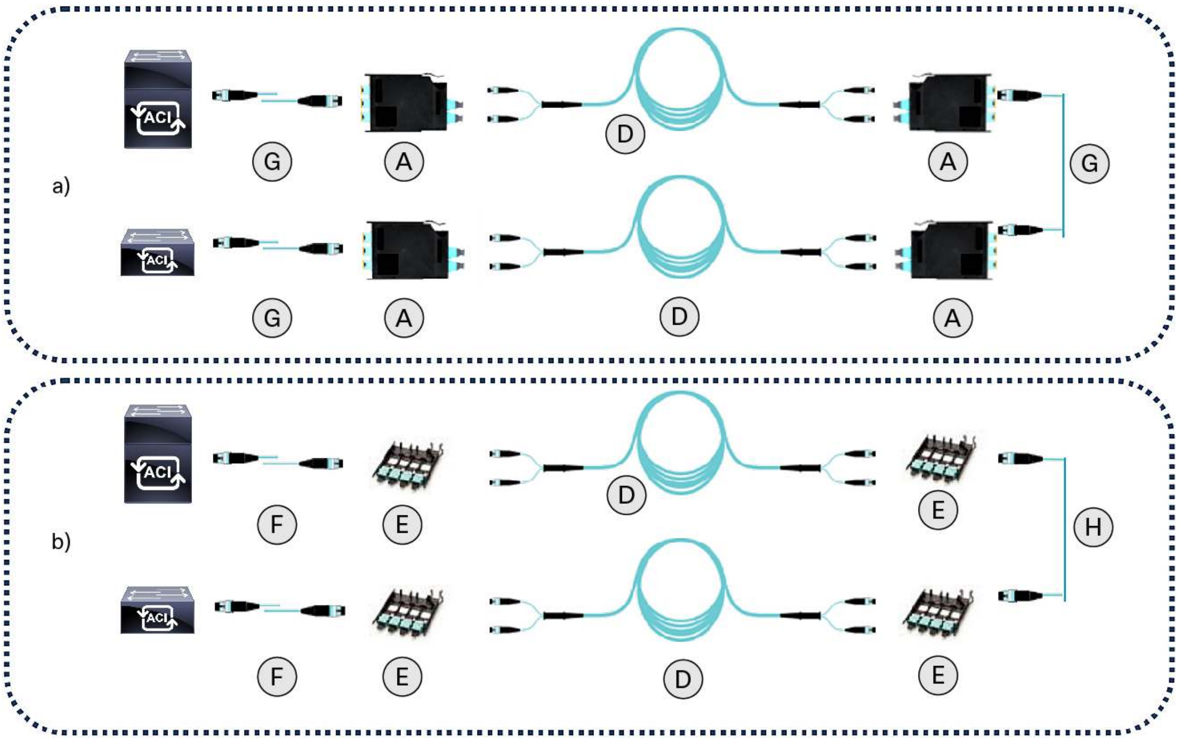

Table 9. Sample Part Numbers for Parallel Optics Cabling Options

Identifier	Part Number	Description
A	ECM-UM24-93-93Q	2x3 conversion module; MTP (pinned) to MTP (pinned) and OM3/4 cable **Note:** A higher-density 4x6 module also is available.
D	G757524TPNDDU100F	24-fiber MTP OM3 trunk cable; 100 ft **Note:** The example shows a 24-fiber trunk cable, but trunk cables can have 12 to 144 fibers at any length.
E	EDGE-CP48-E3	48-fiber (4-port) MTP adapter panel
F	J937512TE8-NB010F	MTP (pinned) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft
G	J757512TE8-NB010F	MTP (pinless) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft

Conclusion

Structured cabling using an MTP cabling infrastructure can be used with current 10 Gigabit Ethernet environments while maintaining investment protection for 40-Gbps environments and beyond. With the new 40 Gigabit Ethernet bidirectional transceivers, no changes to the cabling infrastructure are required when transitioning from 10 to 40 Gigabit Ethernet. Extended 40 Gigabit Ethernet link distances, which match the distances at 10 Gigabit Ethernet, can be achieved by converting to parallel optics transceivers. These transceivers require a change in traditional cabling practices. However, if structured cabling has been implemented with MTP-based trunk cables, then making the conversion is as simple as swapping the patch panels. Thus, the existing MTP-LC modules that were used in the two-fiber serial transmission would be replaced with MTP conversion modules for parallel optics.

New data center switching platforms, such as the Cisco Nexus 9000 Series, are now using the cost-effective, lower-power optics at 40 Gbps to deploy innovative and flexible networking solutions. These solutions allow easy integration into existing environments and deployment of new options regardless of your zone (EoR or MoR) or ToR deployment needs.

For More Information

For more information on Cisco data center switch platforms such as Nexus 9000 Series Switches, visit the website http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html.

For more information on Cisco optical transceiver products, visit the website http://www.cisco.com/en/US/products/hw/modules/ps5455/prod_module_series_home.html.

For additional assistances in designing you cabling infrastructure or additional information on Corning cabling solutions, contact a Corning customer service representative. Likewise, if you would like to read Corning's design guide or request that the support team contact you, visit http://cablesystems.corning.com/1-NX6Cabling.

About Cisco

Cisco is a worldwide leader in networking that transforms the way that people connect, communicate, and collaborate. Information about Cisco can be found at http://www.cisco.com.

About Corning Cable Systems

Corning Cable Systems, part of the Corning Incorporated telecommunications segment, is a leading manufacturer of fiber-optic communications system solutions for voice, data, and video network applications worldwide. For more information, visit the website at http://www.corning.com/cablesystems.

Appendix A: Bills of Materials

Table 10 provides ordering information for Cisco® 40-Gbps bidirectional and parallel optical transceiver products. Table 11 provides ordering information for Corning 40-Gbps bidirectional and parallel optical transceiver products.

Table 10. Cisco Part Numbers

Part Number	Image	Description
QSFP-H40G		QSFP+ Twinax cables
QSFP-40G-SR4		QSFP+ SR4 optics; used with 50 micron MMF
QSFP-40G-CSR4		QSFP+ CSR4 extended reach optics; used with 50-micron MMF
QSFP-40G-SR-BD		QSFP+ SR bidirectional optics; used with 50-micron MMF
SFP-10G-SR **SFP-10G-SR-X**		SFP+ SR optics; used with 50-micron MMF

Table 11. Corning Part Numbers

Identifier	Part Number	Image	Description
A	ECM-UM24-93-93Q		2x3 conversion module; MTP (pinned) to MTP (pinned) OM3/4 cable **Note:** A higher-density 4x6 module is also available.
B	ECM-UM12-05-93T		12-fiber LC duplex to MTP (pinned); OM3
C	797902TD120002M		LC uniboot OM3 jumper; 2m
D	G757524TPNDDU100F		24-fiber MTP OM3 trunk cable; 100 ft **Note:** The example uses a 24-fiber trunk cable, but trunk cables can have 12 to 144 fibers at any length.
E	EDGE-CP48-E3		48-fiber (4-port) MTP adapter panel
F	J937512TE8-NB010F		MTP (pinned) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft
G	J757512TE8-NB010F		MTP (pinless) to MTP (pinless) OM3 jumper with Type-B polarity; 10 ft
H	J939312TE8-NB010F		MTP (pinned) to MTP (pinned) OM3 jumper with Type-B polarity; 10 ft
I	H937524QPHKLZxxxF		2x3 conversion harness assembly; 12-fiber MTP connectors are pinned, and 8-fiber MTP connectors are pinless; 24-inch breakout legs; OM3/4; xxx ft
K	EDGE-04U		Edge 4RU 19-inch rack-mount housing; holds 48 edge modules or panels

Learn more about each of these products in the Corning product catalog; see “Indoor Preterminated Systems” at http://catalog.corning.com/CableSystems/en-US/Default.aspx.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-729908-00 11/13