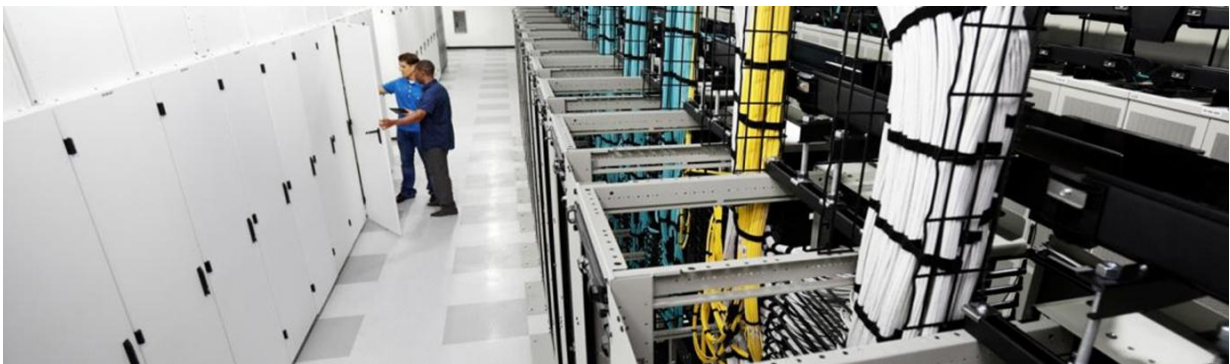# Simplified 40-Gbps Cabling Deployment Solutions with Cisco Nexus 9000 Series Switches

Panduit and Cisco Accelerate Implementation of
Next-Generation Data Center Network Architecture

Contents

What You Will Learn

With the expanding bandwidth needs that virtualization and cloud technologies have applied to the data center network, staying ahead of the curve is important to maintain a competitive advantage. This document introduces products from Cisco and Panduit that address the associated needs. This document also maps the logical designs to the physical layout and displays these products deployed within the data center network. The first section of this document introduces the Cisco Nexus® 9000 Series Switches and typical deployment scenarios and use cases. The second section describes the cabling scenarios and the optics needed for the interworking of the products from both companies. The last section of the document summarizes some Panduit product offerings and next steps.

Introduction

Data centers are undertaking some profound changes due to the ongoing data center consolidation, virtualization and cloud technology, and increasing demand on application workload extensity and agility. These driving forces are changing every facet of the data center, flattening its topology from the traditional core-aggregation-access 3-tier design to the spine-leaf 2-tier architecture, boosting its bandwidth capacity from 1GE at access and 10GE at aggregation to 10GE/25GE at leaf and 40GE/100GE at spine, and morphing its operational model from per-box CLI manual approach to system-level automated process through Restful API. While performance, bandwidth and latency remain as the most important elements in data center networks, management and operational agility and simplicity have elevated to the top of mind of data center architects and operators.

Cisco Nexus 9500, 9300, and 9200 Switches are the next-generation products aiming to take data center networks through all these transitions and beyond. With their unprecedented performance and port density, these switches are well suited in both the traditional 3-tier design and the spine-leaf design providing forwarding capacity for fast increasing server-to-server traffic. They are also equipped with unmatched programmability and automation functionalities to transform the data center network management model.

Cisco Nexus 9000 Series Switches

The Cisco Nexus 9000 Series Switches are the foundation of the Cisco next-generation data center solution. Designed to meet the requirements of mission-critical data centers, these switches deliver exceptional availability, outstanding scalability, and the enhanced Cisco NX-OS Software data center switching feature set.

The Cisco Nexus 9000 Series Switches are the latest extension to the Cisco Nexus Family of modular switches. It encompasses the Nexus 9500 Series modular switches and the Nexus 9300 and 9200 fixed-port switches. The Cisco Nexus 9504, 9508, and 9516 make up the modular 9500 switch platform. The 9332PQ, 9336PQ, 9372PX-E, 9372TX-E, 9396TX, 9396PX, 93120TX, 93128TX, 93108TC-EX, and 93180YC-EX switches are members of the Cisco Nexus 9300 switch platform. Finally, the 9272TX and 92160YC-X comprise the 9200 switch platform. All 9000 switches are high-density 10, 40, and 100 Gigabit Ethernet switches designed for a broad range of environments including physical and virtual computing, Big Data applications, and high-performance computing (HPC).

Cisco Nexus 9500 Modular Switches

The Cisco Nexus 9504 (Figure 1, left switch) is a 7RU switch offering 4 line card slots. The Cisco Nexus 9508 (Figure 1, middle switch) is a 13RU switch offering 8 line cards slots. The Cisco Nexus 9516 (Figure 1, right switch) is a 21RU switch offering 16 line card slots. All of these switches offer a comprehensive selection of modular line cards. The table below displays the port capacities and options for each chassis.

Figure 1: Cisco Nexus 9500 Switch Family

Cisco Nexus 9504   Cisco Nexus 9508   Cisco Nexus 9516
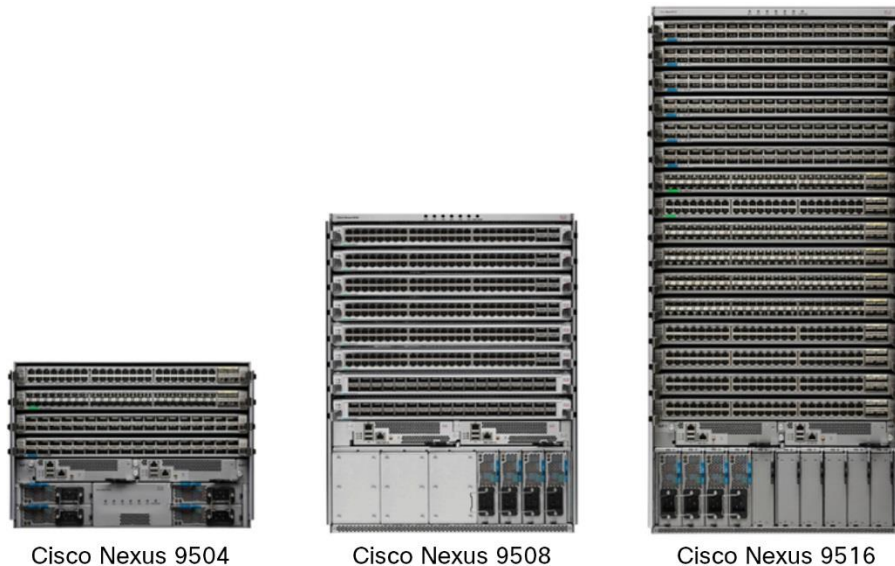
	9504	9508	9516
Nonblocking Performance	15Tbps	30Tbps	60Tbps
10 Gigabit Ethernet Nonblocking Ports	576	1152	2048
40 Gigabit Ethernet Nonblocking Ports	144	2x88	576
1/10GBASE-T Nonblocking Ports	192	384	768
Form Factor	7RU	13RU	21RU
Line Card Slots	4	8	16
Airflow	Front to Back	Front to Back	Front to Back

Cisco Nexus 9500 Line Card Offerings

The Cisco Nexus 9500 platform supports a variety of line card options:

1 and 10 Gigabit Ethernet

The Cisco 9500 platform offers four options for 1 and 10GE line cards. The details of each of the cards are listed in Table 1 below and can be seen in Figure 2.

Table 1: Cisco Nexus 9500 Switch 1/10GE Modules

	9564PX	9464PX	9564TX	9464TX
Usage	10 aggregation, and FEX[*]	10 aggregation and FEX[*]	End of row	End of row
Front-panel ports	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE	48 x 10 GE + 4 x 40 GE
Standalone	Yes	Yes	Yes	Yes
ACI leaf	No	No	No	No
ACI spine	No	No	No	No

	9564PX	9464PX	9564TX	9464TX
Chassis support	4, 8, and 16 slot	4, 8, and 16 slot	4, 8, and 16 slot	4, 8, and 16 slot
100 M/1 GE/10 GE copper ports	NA	NA	48	48
10 GE fiber ports	48	48	NA	NA
40 GE ports	4	4	4	4
Switching performance	Nonblocking	Line rate greater than 200 byte	Nonblocking	Line rate greater than 200 byte
Buffer (MB)	104	12	104	12
Typical power	300	160	450	320
Minimum software version	6.1(2)I2	6.1(2)I2(2a)	6.1(2)I2	6.1(2)I2(2a)
SKU	N9K-X9564PX	N9K-X9464PX	N9K-X9564TX	N9K-X9464TX

Figure 2: Cisco Nexus 9500 Switch 1/10GE Modules

Cisco N9K-X9464PX

Cisco N9K-X9464TX

Cisco N9K-X9564PX

Cisco N9K-X9564TX

40 Gigabit Ethernet

The Cisco 9500 Platform offers three options for 40GE line cards. The details of each of the cards are listed in Table 2 below and can be seen in Figure 3.

Table 2: Cisco Nexus 9500 Switch 40GE Modules

	9636PQ	9536PQ	9432PQ
Usage	40 GE high-performance aggregation and FEX[*]	40 GE aggregation, ACI leaf, and FEX[*]	40 GE aggregation and FEX[*]
Front-panel ports	36 x 40 GE	36 x 40 GE	32 x 40 GE
Standalone	Yes	Yes	Yes
ACI leaf	No	No	No
ACI spine	No	No	No
Chassis support	4 and 8 slot	4, 8, and 16 slot	4, 8, and 16 slot
40 GE ports	36	36	32
Switching performance	Nonblocking	1.5:1 oversubscribed	Line rate more than 200 byte
Buffer (MB)	36	104	24
Typical power (watts)	260	360	240
Minimum software version	6.1(2)I1(1)	6.1(2)I2(2a)	6.1(2)I2(2a)
SKUs	N9K-X9636PQ	N9K-X9536PQ	N9K-X9432PQ

Figure 3: Cisco Nexus 9500 Switch 40GE Modules

Cisco N9K-X9432PQ   Cisco N9K-X9536PQ   Cisco N9K-X9636PQ

100 Gigabit Ethernet

The Cisco 9500 Platform offers two options for 100GE line cards. The details of each of the cards are listed in Table 3 below and can be seen in Figure 4.

Table 3: Cisco Nexus 9500 Switch 100GE Modules

	X9408PC-CFP2	X9432C-S
Usage	Backbone and router interconnect	Aggregation and spine
Front panel ports	8 x 100 GE	32 x 100 GE
Multispeed ports	N/A	4 x 1 GE, 4 x 10 GE, 4 x 25 GE, 1 x 40 GE, 2 x 50 GE
Standalone	Yes	Yes
ACI leaf	No	No
ACI spine	No	No
Chassis support	4-, 8-, and 16-slot	4-[*], 8-, and 16-slot[*]
Switching performance	Nonblocking	Line rate greater than 250 bytes
Interface type	CFP2	QSFP28
Buffer (MB)	24	32
Typical power	410	340
SKU	N9K-X9408PC-CFP2	N9K-X9432C-S

Figure 4: Cisco Nexus 9500 Switch 100GE Modules

Cisco N9K-X9408PC-CFP2   Cisco N9K-X9432C-S

Cisco Nexus 9300 Fixed-Port Switches

The Cisco Nexus 9300 Series Switches are fixed switches built to provide 1/10Gbps and 40Gbps server access connectivity. The 1RU options include the Nexus 9332PQ (Figure 5, left switch), 9372TX-E (Figure 5, middle switch), and 9372PX-E (Figure 5, right switch). The 2RU options include the Nexus 9396TX (Figure 6, left switch), 9396PX (Figure 6, middle switch), and 93120TX (Figure 6, right switch). The Cisco Nexus 93128TX (Figure 7) is a 3RU fixed switch with 96 ports of 10GBase-T and 8 ports of 40Gbps (QSFP+). All 9300 switches are well suited for data center top-of-rack (ToR), end-of-row (EoR), and collapsed aggregation and access while deployed with Nexus 2000 Fabric Extenders.
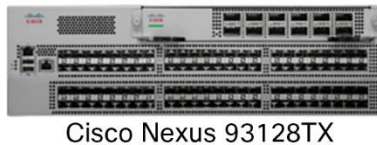
Figure 5: Cisco Nexus 9300 Switch 1RU Switches

Cisco Nexus 9332PQ   Cisco Nexus 9372TX-E   Cisco Nexus 9372PX-E

Figure 6: Cisco Nexus 9300 Switch 2RU Switches

Cisco Nexus 9396TX   Cisco Nexus 9396PX   Cisco Nexus 93120TX

Figure 7: Cisco Nexus 9300 Switch 3RU Switches

Cisco Nexus 93128TX

Cisco Nexus 9300-EX Fixed-Port Switches

The Cisco Nexus 93180YC-EX Switch (Figure 8) is a 1-rack-unit (1RU) switch with latency of less than 1 microsecond that supports 3.6 terabits per second (Tbps) of bandwidth and over 2.8 billion packets per second (bpps). The 48 downlink ports on the 93180YC-EX can be configured to work as 1-, 10-, or 25-Gbps ports, offering deployment flexibility and investment protection. The uplink can support up to six 40- and 100-Gbps ports, or a combination of 10-, 25-, 40-, 50-, and 100-Gbps connectivity, offering flexible migration options. The Cisco Nexus 93108TC-EX Switch (Figure 9) is a 1RU switch that supports 2.16 Tbps of bandwidth and over 1.7 bpps. The 48 10GBASE-T downlink ports on the 93108TC-EX can be configured to work as 100-Mbps, 1 Gbps, or 10-Gbps ports. The uplink can support up to six 40- and 100-Gbps ports, or a combination of 10-, 25-, 40, 50-, and 100-Gbps connectivity, offering flexible migration options.

Figure 8: Cisco Nexus 93180YC-EX Switch

Figure 9: Cisco Nexus 93108TC-EX Switch

Cisco Nexus 9200 Fixed-Port Switches

The Cisco Nexus 9200 Series Switches are fixed switches built to provide 1/10Gbps and 40Gbps server access connectivity. The Cisco Nexus 92160YC-X (Figure 10) is a 1RU fixed switch offering 64 ports of 10 Gbps or 25Gbps (SFP+), and 6 ports of 40 Gbps or 4 ports of 100 Gbps (QSFP28) with a nonblocking switch architecture. The Cisco Nexus 9272Q (Figure 11) is a 2RU fixed switch with up to 144 ports of 10GBase-T and 72 ports of 40Gbps (QSFP+). All Nexus 9200 switches are well suited for data center ToR, EoR, and collapsed aggregation and access while deployed with Nexus 2000 Fabric Extenders.

**Figure 10:** Cisco Nexus 92160YC-X 1RU Switch

Figure 11: Cisco Nexus 9272Q 2RU Switch

Cisco Nexus 9000 Series Switches Deployment Scenarios

Integrated line-rate Layer-2 and Layer-3 along with a comprehensive feature set make the Cisco Nexus 9000 Series Switches excellent for deployments in enterprise, service provider, and commercial environments.

The Cisco Nexus 9500Switches can be deployed with Cisco Nexus Fabric Extenders in both EoR and MoR designs. For an EoR or MoR design, servers can use standard RJ-45 cables to ToR patch panels in the rack. The bundles of copper cables from each rack are routed to the EoR or MoR network racks and to the appropriate Cisco Nexus Switch or Cisco Nexus Fabric Extender with a Cisco Nexus parent switch as shown in Figure 12. Connectivity from the Cisco Nexus 9000 Series Switches to the aggregation layer Cisco Nexus 9000 Series Switches can be performed by fiber optic cabling (OM3 or OM4) or copper cabling (RJ-45 or Direct-Attached Copper Cabling).
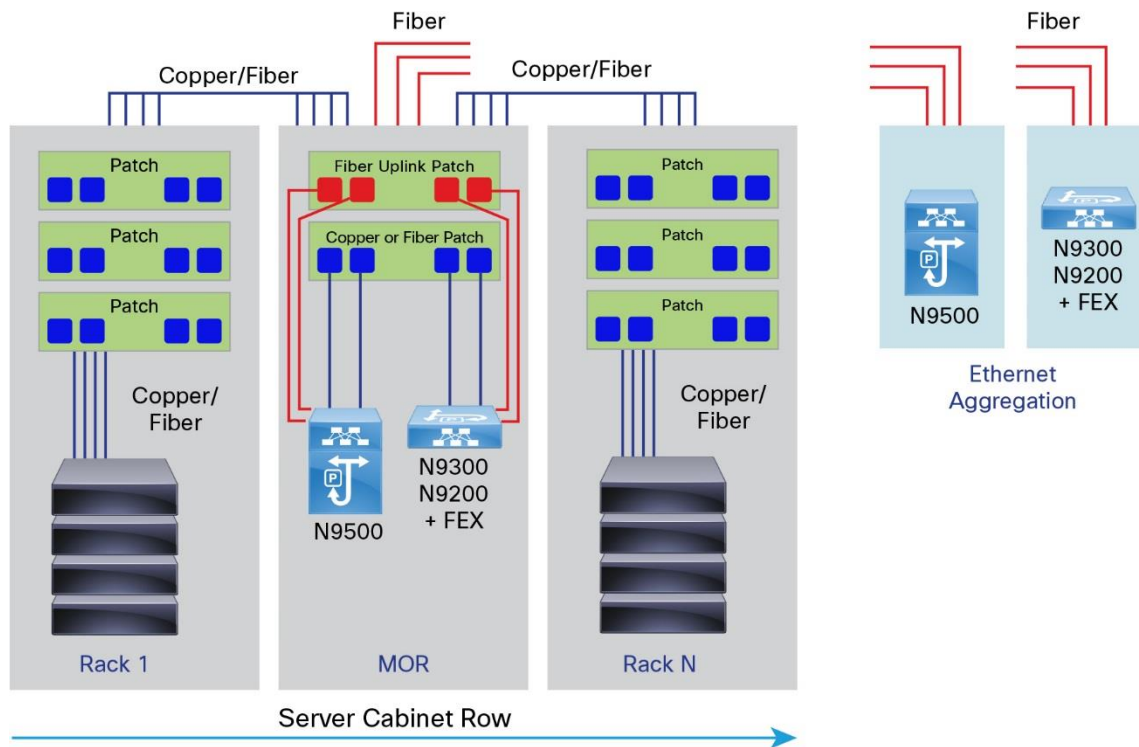
Figure 12: EoR Design Using Cisco Nexus 9000 Series Switches
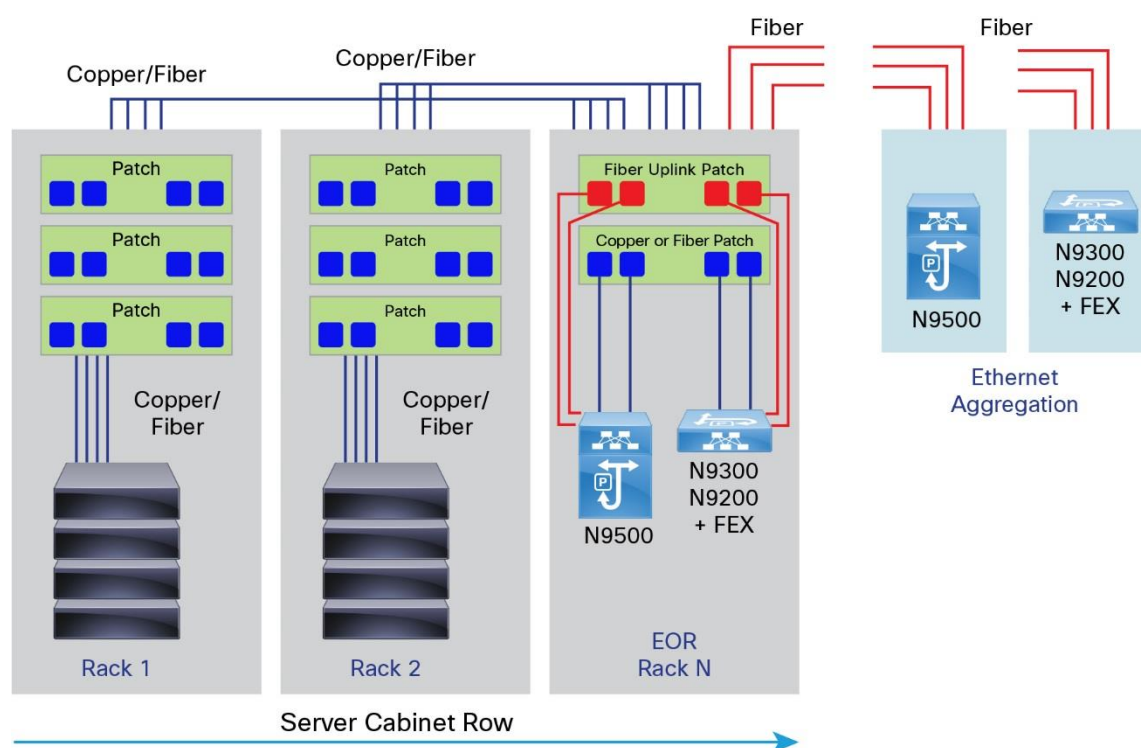
In MoR deployments, instead of long cable runs, the copper cable is routed from each server rack to a pair of racks positioned next to each other in the middle of the row. This approach reduces the extreme cable lengths from the far-end server cabinets using an EoR design. In both the EoR and MoR designs, if the trunks and ToR panels are fiber based, then it is possible to use fiber-based fabric extenders or connect directly to Cisco Nexus 9000 Series Switches (Figure 13).
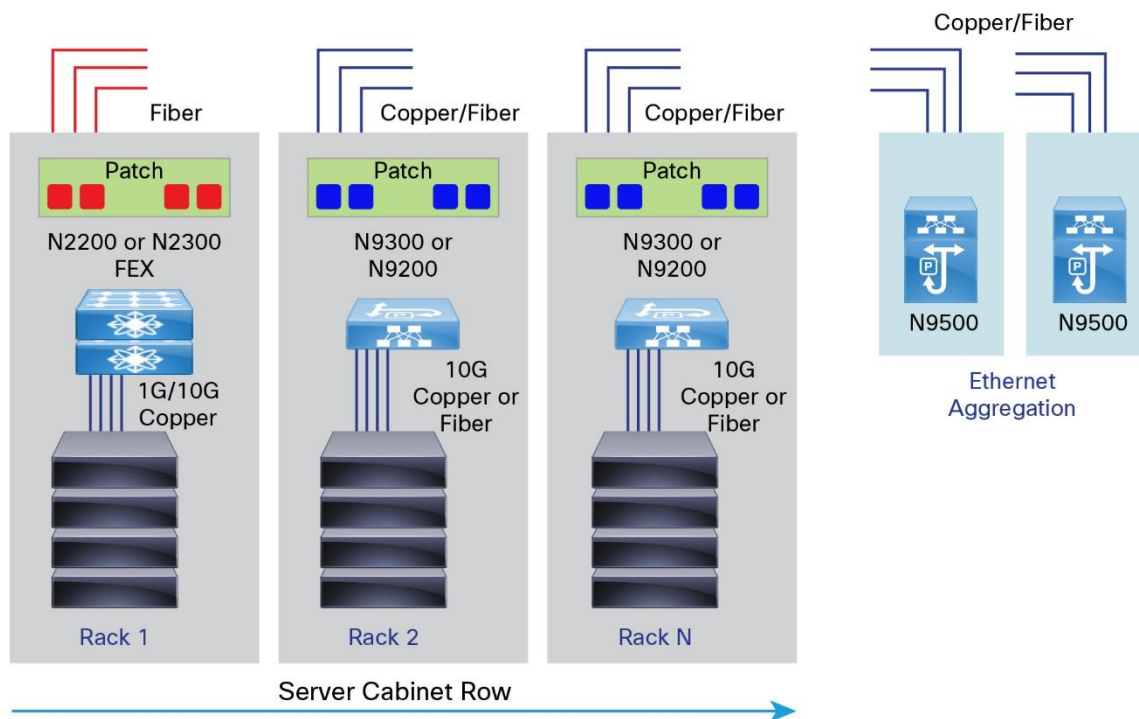
Figure 13: MoR Design Using Cisco Nexus 9000 Series Switches

With an EoR or MoR design using the Cisco Nexus 9000 Series Switches with fabric extenders, there are fewer switches to manage, reducing capital expenditures (CapEx). The Cisco Nexus 2000 Series Fabric Extenders are logically managed as part of the Cisco Nexus 9000 Series Switches, so you get the additional benefit of lower operating expenses (OpEx) because there are fewer management points.

In ToR designs, various devices can be used such as the Cisco Nexus 3000 Series Switches, Cisco Nexus 2000 Series Fabric Extenders, and any Cisco Nexus 9300 Fixed-Port Switches. In ToR designs, the servers in the computer rack connect directly to either a Cisco Nexus 9300 Fixed-Port Switch or a Cisco Nexus 2000 Series Fabric Extender in the same rack. All cabling for servers stays within the racks, as short patch cables, Twinax cables, or fiber from the server to the ToR network device. From each rack, a fiber would run directly to an aggregation switch. For shorter distances not exceeding 10 meters, passive or active Twinax cabling can also be used between the ToR and aggregation. A ToR design has the benefit of less bulky and less expensive copper cabling between racks and throughout the data center. With less copper cabling, there is less airflow obstruction. This allows the network to function with fewer racks and there is less infrastructure needed for patching and cable management (Figure 14).

Figure 14: ToR Design using Nexus 9000 Series Switches

For the ToR design in Figure 14, the Cisco Nexus 2200/2300Switches will connect to the parent switch in the aggregation layer (Cisco Nexus 9500 Modular Switch).

Cisco Nexus 9000 Series Switches Use Case

The Cisco Nexus 9000 Series Switches are a strong fit for the data center use cases described here.

Traditional 3-Tier Architecture

In a traditional data center network, it is common that the network is deployed in the 3-tier architecture, including data center core, data center aggregation and data center access as shown in Figure 15. Nexus 9000 Series Switches have a natural fit into this design with the Nexus 9500Switches at the aggregation layer and Nexus 9500, 9300, or 9200 Series Switches at the access layer.

Figure 15: Traditional 3-Tier Network Architecture
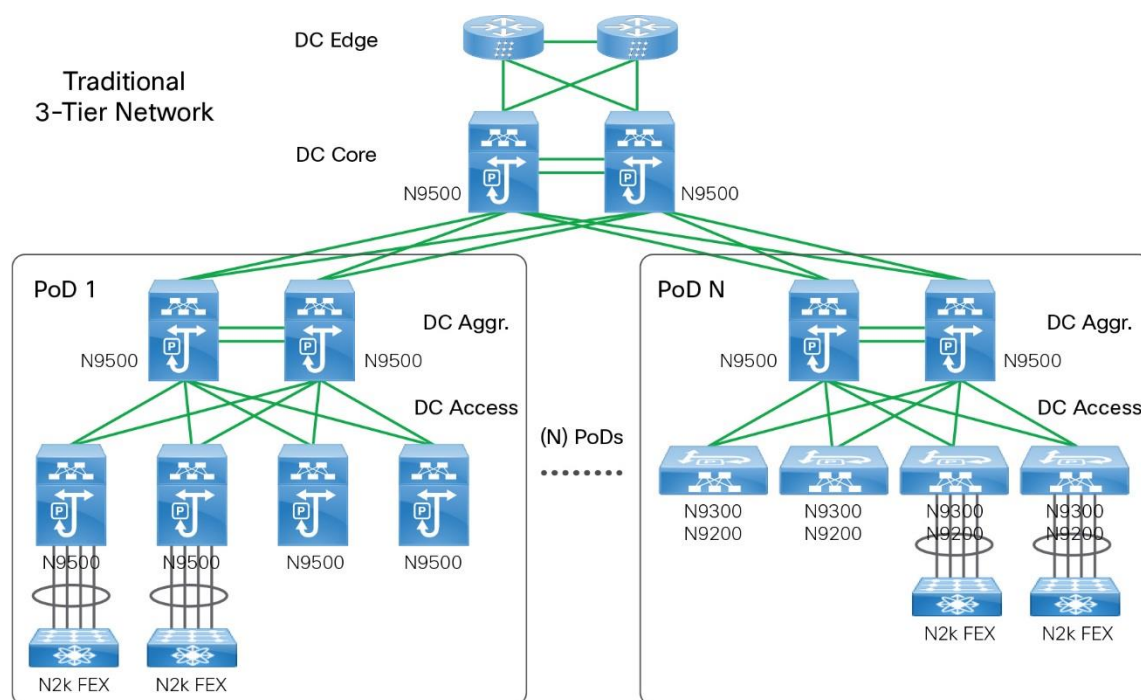
Large-Scale Fabric (Layer 2 and 3) with Leaf-and-Spine Architecture

To meet the growing demands for scale, performance, and flexibility in today's data centers, the trend in network design is to flatten the architecture using high-density and low-latency switches as leaves and spines in a Layer-2 or Layer-3 network. A network built using a leaf-and-spine architecture allows connections scaling from hundreds to 10,000 or more servers with high east-west traffic flows. The leaf-and-spine architecture helps ensure low latency with cut-through forwarding as packets traverse the network using a low hop count. The spine switch creates a nonblocking, low-latency fabric, forwarding packets between leaves. The leaf switches provide connectivity to servers. Use of a CLOS architecture helps ensure the highest possible network availability with little impact on customer traffic if a failure occurs. In the topology in Figure 16, the Cisco Nexus 9000 Series Switches are deployed in a leaf-and-spine architecture with Nexus 9500 Switches as spines and Nexus 9300, 9200 or 9500Switches as leaf switches. The Nexus 2000 Fabric Extender can be attached to either Nexus 9300, 9200 or Nexus 9500Switches to provide cost-effective server connectivity, especially 1Gbps server connectivity.

Figure 16: Large-Scale Fabric: Leaf-and-Spine Architecture
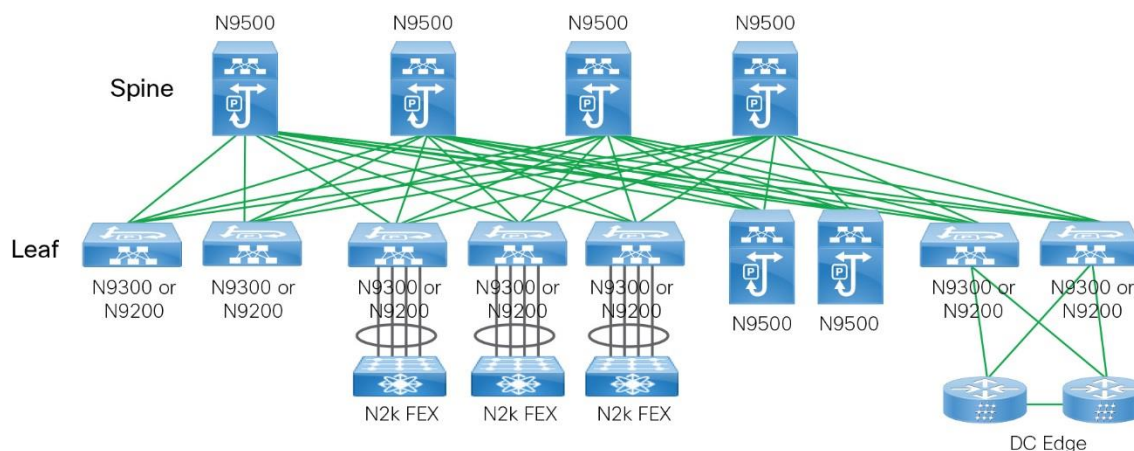
Fabric Extender Architecture: High-Density Fabric Extender Aggregator

Cisco FEX Technology is based on the emerging IEEE 802.1br standard. Designing the network using Cisco FEX Technology provides flexibility, reduces cabling infrastructure, and provides a single point of management, helping customers scale their networks. With its high 10- and 40-Gbps port density, the Cisco Nexus 9000 Series Switches can be used as a high-density fabric extender aggregation platform for the Cisco Nexus 2200 platform products (Figure 17).

Figure 17: Fabric Extender Aggregation
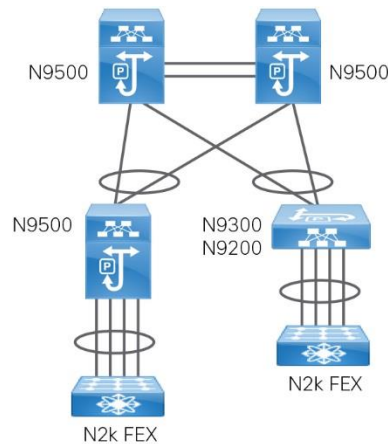
The Cisco Nexus 2000 Series Fabric Extenders can be physically located at the top of the rack, and the Cisco Nexus 9000 Series Switch can reside in the middle of the row or at the end of the row. This positioning takes advantage of both EoR network consolidation and ToR cabling savings. The Cisco Nexus 9000 Series Switches are compatible with integrated transceivers and Twinax cabling solutions, which deliver cost-effective connectivity for 1 or 10 Gigabit Ethernet to servers at the rack level, eliminating the need for expensive optical transceivers.

Cisco and Panduit® Wiring Scenario

Using the three topologies discussed in the previous section, the scenarios in this section present the high-level cabling and optics requirements for the Cisco Nexus 9000 Series Switches. Typical wiring solutions and scenarios are described here and include products from Cisco and Panduit.

The wiring scenarios described here can be used in a Cisco Nexus 9000 Series Switch deployment.

PanMPO™ Connector

The PanMPO Connector allows end users to change the polarity and gender of the MPO fiber connector in the field, with no risk of damaging the internal fiber ribbon or the connector's endface. Using reference patch cords terminated with the PanMPO Connector simplifies the process of both testing and certifying existing cable plant and new installations, which helps ease the migration from 10Gb/s Ethernet to 40Gb/s Ethernet, while keeping the installation compliant with standards and industry best practices. Figure 18 shows PanMPO Connector options.
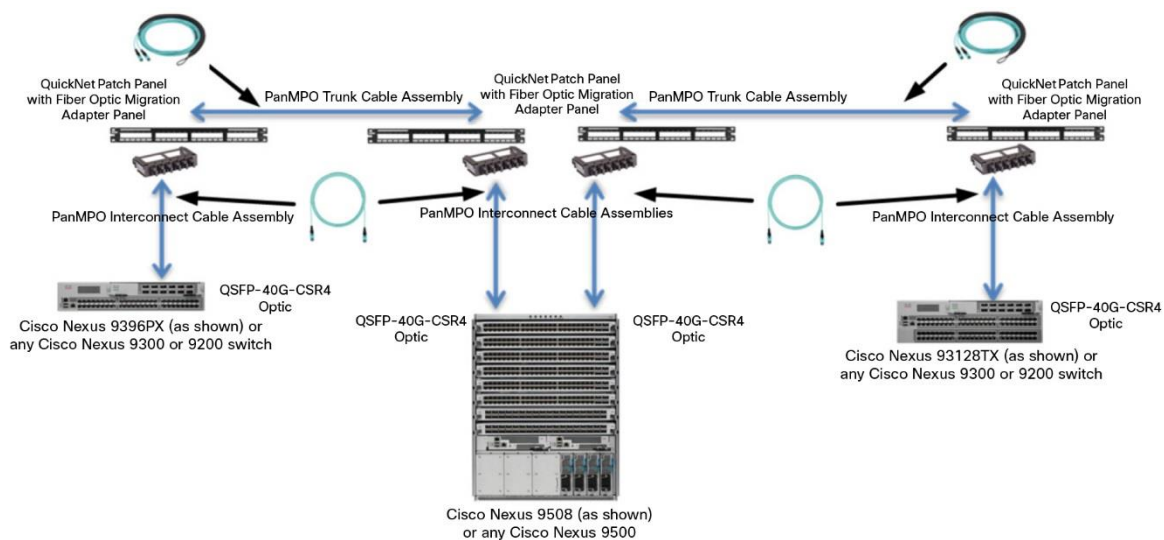
Figure 18: PanMPO Connector Options
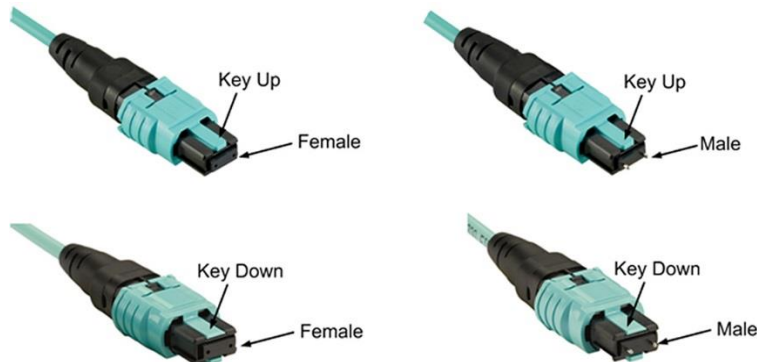
40GE QSFP SR4/CSR4 Optics Cabling Options

Figure 19 shows this scenario.

Figure 19: Cisco Nexus 9396 and 93128 Switches to Cisco Nexus 9508 PanMPO to PanMPO Multimode Cabling Components

For distances less than 400 meters, the PanMPO multimode fiber cabling is the preferred cabling method. Figure 19 shows a wiring option for connecting a Cisco Nexus 9396 to a Cisco Nexus 9508 and to a Cisco Nexus 93128 Switch using PanMPO interconnect cable assemblies. With Cisco QSFP-40G-SR4optics, this scenario is capable of distances of up to 100 meters on OM3 fiber or 150 meters with OM4 fiber. The QSFP-40G-CSR4 optics can be used for distances of up to 300 meters with OM3 fiber or 400 meters with OM4 fiber. With Panduit Signature Core™ fiber cabling, this scenario can be used for distances up to 550 meters.

Panduit's Signature Core Fiber Optic Cabling System is a standards compliant advanced fiber optic cabling system that counterbalances the dispersive effects of both modal dispersion and chromatic dispersion, therefore minimizing total dispersion. The system delivers signal integrity far beyond the requirements for 10/40/100 Gb/s Ethernet, and 8 and 16 Gb/s Fiber Channel, providing the ultimate in design flexibility which allows implementation of complex Data Center architectures, and verified optical performance. It improves MMF performance by increasing modal bandwidth and ensures consistent performance and reliability of critical systems, delivering the ability to deploy more connectors in the channel, which simplifies moves, adds and changes.

From the Cisco Nexus 9396 port, the QSFP-40G-SR4 or QSFP-40G-CSR4 optics connects to an PanMPO fiber interconnect cable to the front of the fiber optic adapter panel. The fiber optic adapter panel fits into the QuickNet™ patch panel. The PanMPO fiber trunk cable plugs into the back of the fiber optic adapter panel using the port that corresponds to the port that the interconnect cable used in the front. On the remote end, the PanMPO fiber trunk cable plugs into the back of the fiber optic adapter panel. The PanMPO fiber interconnect cable plugs into the corresponding port on the front of the panel on one side and into the optics with the switch on the other end.

The PanMPO interconnect fiber assembly used in the scenario is a Female PanMPO-to-Female PanMPO cable, and the PanMPO trunk cable assembly is a Male PanMPO-to-Male PanMPO cable. With this cabling method, the interconnect (or intracabinet) cables will always be Female to Female, because MPO optics in a switch or router are always male connectors per the IEEE standards and will always accept Female PanMPO connectors. The trunk cabling will then always be Male PanMPO-to-Male PanMPO trunk cables. The use of this cabling methodology throughout the data center allows a simplified cabling scheme and quick identification of interconnect and trunk cabling.

In addition to using the Female PanMPO for interconnect cables and the Male PanMPO for trunk cables, users should use Method B polarity (key up to key up) cabling for both the interconnect and trunk cables. This approach standardizes the type of cable used throughout the data center, eliminating the need for several different interconnect cable types, as would be necessary if other polarity methods in the infrastructure are used.

Table 4 shows the Cisco and Panduit optics and cabling options with corresponding part numbers for a QSFP+-to-QSFP+ multimode connection. Figure 20 shows a PanMPO cabling scenario.

Table 4: Cisco Nexus 9000 Series Switches PanMPO to PanMPO Multimode QSFP+ Optics and Cabling Products

Cisco Product Description	Cisco Part Number
40GBASE SR4 QSFP, 100m with OM3 or 150m with OM4	QSFP-40G-SR4
40GBASE Extended SR4 QSFP, 300m MMF with OM3 or 400m with OM4	QSFP-40G-CSR4
Panduit Product Description	**Panduit Part Number**
QuickNet 24-Port Patch Panel	QPP24BL
QuickNet Fiber Optic Migration Adapter Panel	FQMAP65CG
QuickNet PanMPO Interconnect Cable Assembly (OM4) (Configured with Female to Female, Method B[*])	FZTRP7N7NBNF[**]
QuickNet PanMPO Interconnect Cable Assembly (OM3) (Configured with Female to Female, Method B[*])	FXTRP7N7NBNF[**]
QuickNet PanMPO Trunk Cable Assembly (OM4) (Configured with Male to Male, Method B[*])	FZTYP8E8EBAF[**]
QuickNet PanMPO Trunk Cable Assembly (OM3) (Configured with Male to Male, Method B[*])	FXTYP8E8EBAF[**]

[*] PanMPO cable assemblies can change polarity and gender as needed, but should be ordered with the polarity and gender configuration required to accelerate installation.

[**] Insert distance for correct Panduit product number (for example, a 3-foot cable is 003 or FXTYP8E8EBAF003).

Figure 20: Cisco Nexus 9000 Series Switch PANMPO to PANMPO Cabling Diagram
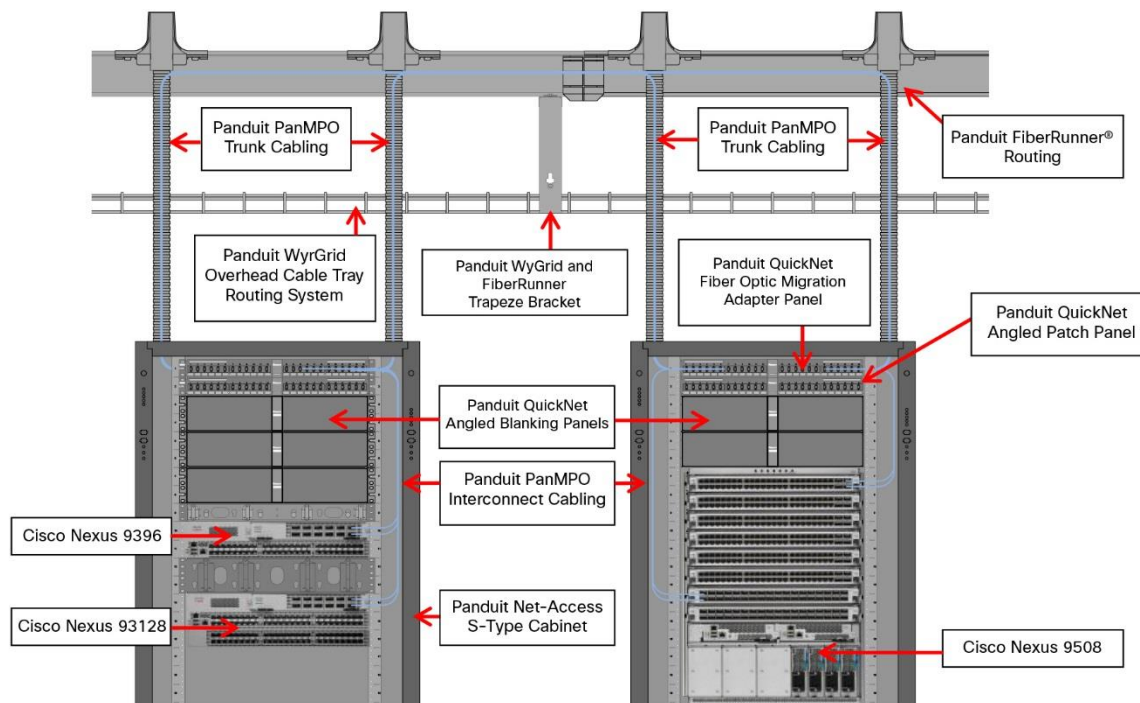
QSFP+-to-SFP+ Configuration with 4 x 10 Gigabit Ethernet

The Cisco Nexus 9500 Series Switch interfaces can operate in 40-Gbps or 4 x 10 Gigabit Ethernet mode. In 40-Gbps mode, each physical port is a single 40-Gbps port with the capabilities to support a 40-Gbps flow. If the interface is logically configured as a 4 x 10 Gigabit Ethernet interface, each port becomes four 10-Gbps ports. The use of Copper Twinax or fiber hydras or breakout cables allows the ports to connect directly to SFP+ ports. When the ports are in 4 x 10 Gigabit Ethernet mode, the Cisco Nexus 9508 switch has a total capacity of 1,152 10-Gbps interfaces per switch.

Through the use of a PanMPO to LC breakout cable, the Cisco Nexus 9500 Series Switches can connect to 10Gbps switch platforms, such as Nexus 3000 Series Switches, or directly to the Cisco Nexus 2000 Series Fabric Extenders. The cable required here would be PanMPO on one end, connecting to the QSFP+ optics in a Cisco Nexus 9500 series switch. It then would break out into four individual fiber pairs links on the other end. These breakout cables terminate with LC duplex connectors.

Cisco Nexus 9508 to Cisco Nexus 3000 Series Switches or the Cisco Nexus 2200 Series Switch using QSFP-to-SFP+ connectivity with breakout cables has the optics options shown here.

PANMPO to (4) LC Hydra

Figure 21 shows this scenario.

Figure 21: Cisco Nexus 9508 Switch PanMPO to LC Duplex Hydra Cabling Components
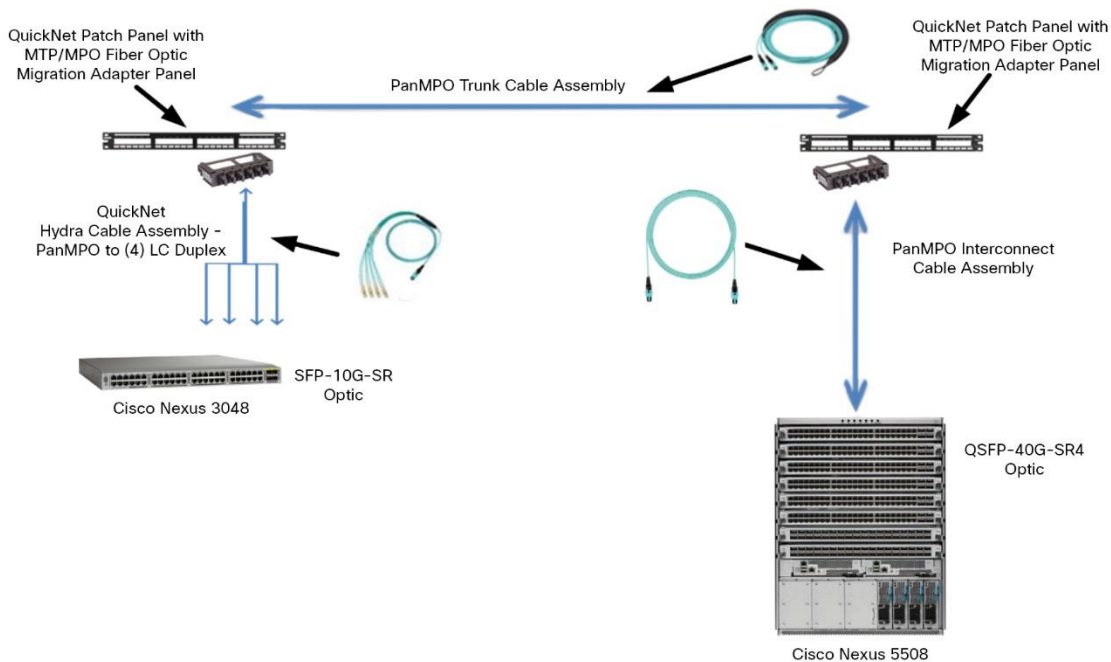
The scenario in Figure 21 shows a wiring option connecting a Cisco Nexus 9508 switch to a Cisco Nexus 3048 switch using a PanMPO 4-to-1 conversion harness cable assembly. With the QSFP-40G-SR4 optics, this scenario is capable of distances of up to 100 meters with OM3 fiber. With the QSFP-40G-CSR4 optics, this scenario is capable of distances of up to 300 meters with OM3 fiber. In this scenario, the Cisco Nexus 9508 has the QSFP-40G-SR4 optics connector, which is connected to a Female-to-Female PanMPO interconnect cable assembly. The Female-to-Female PanMPO interconnect cable assembly connects to the front of the QuickNet patch panel with the PanMPO fiber optic migration adapter panel. On the back side of the PanMPO fiber optic migration adapter panel, the Male-to-Male PanMPO trunk cable assembly connects to the panel in the corresponding port. On the remote side, the Male-to-Male PanMPO trunk cable assembly connects to the back of the panel, and the Female PanMPO 4-to-1 conversion harness cable assembly connects to the front of the panel. The Female PanMPO 4-to-1 conversion harness cable assembly connects from the front of the panel to four corresponding SFP+ ports with SFP-10G-SR optics in the Cisco Nexus 3048 switch using LC duplex plugs.

Table 5 shows the Cisco and Panduit optics and cabling options with corresponding part numbers for a PanMPO 4-to-1 conversion harness cable assembly multimode connection. Figure 22 shows a PanMPO 4-to-1 conversion harness cable assembly cabling scenario.

Table 5: Cisco Nexus 9000 PanMPO 4-to-1 Conversion Harness Optics and Cabling Products

| Cisco Product Description	Cisco Part Number
10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector	SFP-10G-SR
40GBASE SR4 QSFP, 100m with OM3 or 150m with OM4	QSFP-40G-SR4
40GBASE Extended SR4 QSFP, 300m MMF with OM3 or 400m with OM4	QSFP-40G-CSR4

Panduit Product Description	Panduit Part Number
QuickNet 24-Port Patch Panel	QPP24BL
QuickNet Fiber Optic Migration Adapter Panel	FQMAP65CG
QuickNet PanMPO Interconnect Cable Assembly (OM3) (Configured with Female to Female, Method B*)**	FXTRP7N7NBNF**
QuickNet PanMPO Trunk Cable Assembly (OM3) (Configured with Male to Male, Method B*)**	FXTYP8E8EBAF**
QuickNet PanMPO 4-to-1 Conversion Harness Cable Assembly (OM3) (Female MPO to 4 Duplex LC)	FX8RP7NUSQNF**

* PanMPO cable assemblies can change polarity and gender as needed, but should be ordered with the polarity and gender configuration required to accelerate installation.

** Insert distance for correct Panduit product number (for example, a 3 foot cable is 003 or FX8RP7NUSQNF003).

Figure 22 shows a Cisco Nexus 9508 to Cisco Nexus 3048 using PanMPO 4-to-1 conversion harness cable assembly cabling.

Figure 22: Cisco Nexus 9508 PanMPO 4-to-1 Conversion Harness Cable Assembly Cabling Diagram
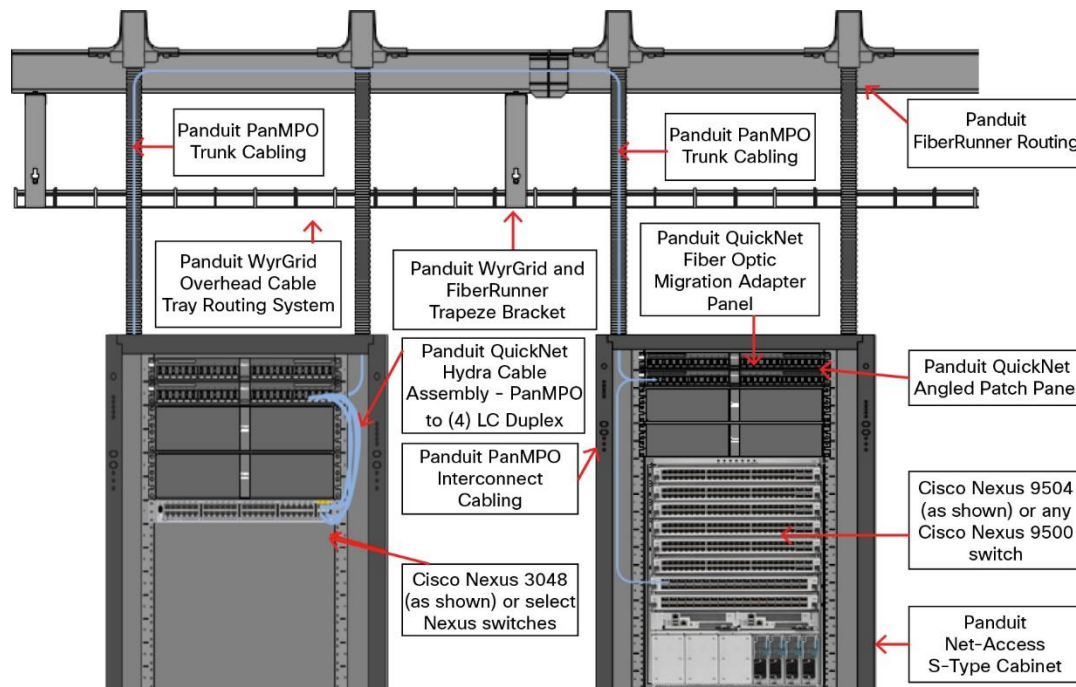
Cisco 40GE QSFP BiDi Optics Cabling Options
Challenges with Existing 40GE Transceivers

Existing short reach transceivers for 40GE in QSFP form factor, such as QSFP SR4 and QSFP CSR4, feature independent transmitter and receiver sections, each with four fiber strands in parallel. For a duplex 40G connection, eight fiber strands are required. Both QSFP SR4 and QSFP CRS4 use the MPO 12-fiber connector. As a result, four fiber strands in each connection get wasted.

Figure 23 depicts the concept of existing short reach 40GE QSFP solution.

Figure 23: Concept of Existing 40GE Transceivers

With existing QSFP transceivers, each direct connection between two devices requires a MPO-to-MPO 12-fiber cable. In the case of structured cabling with patch panels and fiber trunks, a 40GE connection needs MPO-to-MPO fibers between devices and patch panels and four duplex MMF fibers in the fiber trunk.

In most of today's data center networks, the aggregation fiber infrastructure is built for 10GE connectivity that either supports direct connections between devices over LC-to-LC fibers or uses LC-to-LC fibers to attach devices to patch panels and provides one duplex MMF fiber in the fiber trunk for each 10GE connection.

40GE connectivity will not re-use the direct connecting LC-to-LC fibers. It also requires immense expansion to the fiber trunks to fulfill the 4 fibers/40GE-connection requirement. These differences make it expensive for customers to migrate from 10GE to 40GE.

Cisco 40GE QSFP BiDi Transceiver

The QSFP BiDi transceiver addresses the previously mentioned fiber infrastructure challenges by providing the capability of transmitting full duplex 40GE over one duplex MMF fiber with LC connectors. Specifically, the QSFP BiDi transceiver allows for 40GE connectivity to re-use the existing direct connecting 10GE fibers and the existing fiber trunk without additional fibers.

QSFP Bi-Directional (BiDi) is short reach optical transceiver that delivers 40GE over an OM3 or OM4 multimode fiber. It is MSA compliant. Figure 24 shows a QSFP BiDi transceiver.

Figure 24: QSFP BiDi Transceiver (QSFP-40G-SR-BD)

A QSFP BiDi transceiver has two channels of 20G, each transmitted and received simultaneously on two wavelengths over a single multimode fiber strand. The result is an aggregated duplex 40GE link over a duplex multimode fiber. QSFP BiDi uses duplex LC connectors. The connection can reach 100 meters on a OM3 multimode fiber or 125 meters on a OM4 multimode fiber. The Cisco PID for QSFP BiDi transceiver is QSFP-40G-SR-BD. Figure 25 depict the concept of the QSFP BiDi transceiver.

Figure 25: Concept of QSFP BiDi Transceiver
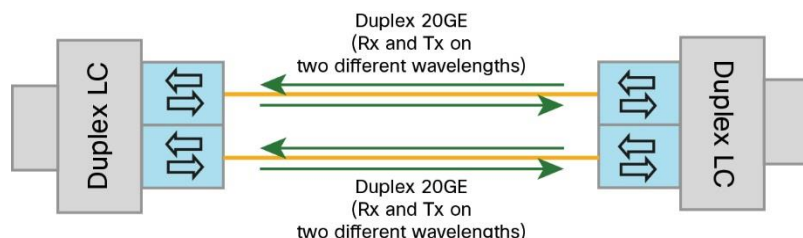
Figure 26: Cisco Nexus 9000 Series BiDi Optics Cabling Components
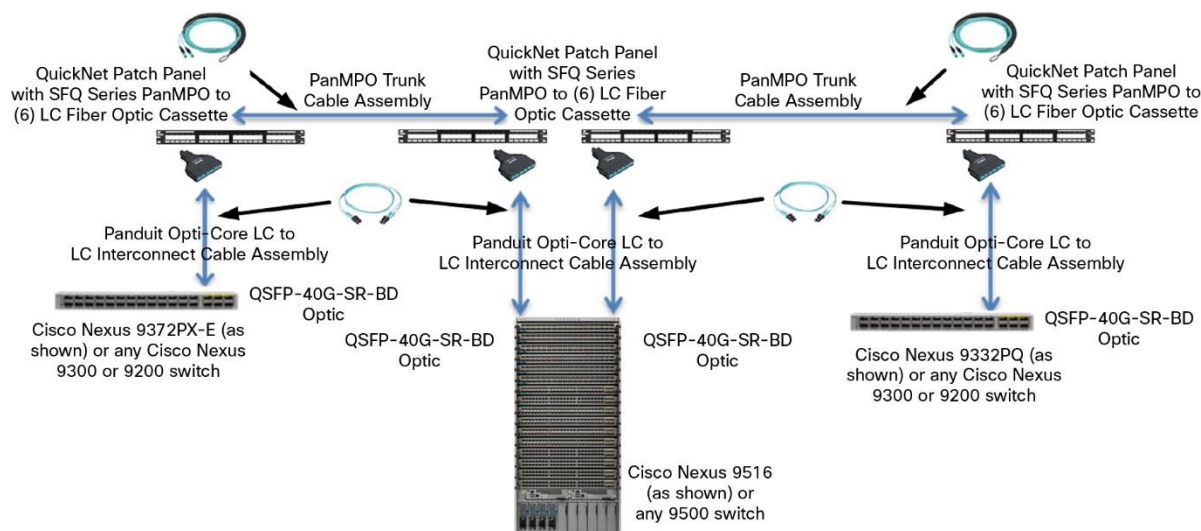

Figure 26 shows a wiring option for connecting a Cisco Nexus 9372PX-E switch to a Cisco Nexus 9516 switch and to a Cisco Nexus 9332PQ switch using PanMPO interconnect cable assemblies.
With Cisco QSFP-40G-SR-BD optics, this scenario is capable of distances of up to 100 meters on OM3 fiber or 125 meters with OM4 fiber.

From the Cisco Nexus 9372PX-E port, the QSFP-40G-SR-BD optic connects to a LC Duplex fiber interconnect cable to the front of the LC fiber optic cassette. The LC fiber optic cassette fits into the QuickNet patch panel. The PanMPO fiber trunk cable plugs into the back of the LC fiber optic cassette. On the remote end, the PanMPO fiber trunk cable plugs into the back of the LC fiber optic cassette. The LC Duplex fiber interconnect cable plugs into the front of the LC fiber optic cassette and into the optics with the switch on the other end.

For a Greenfield Cisco BiDi deployment, Panduit® suggests using Method B polarity, which allows for simplified installation. It requires two different types of cassette to be ordered and installed (B1 & B2), but does not require the LC pairs on one side of the connection to be flipped as would be required in Method A polarity. For ease of installation, it is recommended that you standardize the placement of B1 & B2 cassettes. For example, B1 cassettes are installed in spine or network switch cabinets and B2 cassettes are installed in leaf or server side switch cabinets. Unlike the other scenarios, the PanMPO trunk should be ordered as a Method B Female PanMPO to Female PanMPO, because the cassettes are male PanMPO.

For a Brownfield 10G network upgrading to a 40G Cisco BiDi deployment, customers will only need to insure that each existing 10G link falls within the reach of the 40G Cisco BiDi optic for the installed fiber type (OM3 or OM4). Otherwise no other cabling changes would be required. If the connection was operational at 10G, it will operate at 40G.

Table 6 shows the Cisco and Panduit optics and cabling options with corresponding part numbers for a BiDi multimode connection. Figure 27 shows a BiDi cabling scenario.

Table 6:   Cisco Nexus 9000 Series Switches BiDi Optics and Cabling Products

Cisco Product Description	Cisco Part Number
40GBASE SR-BD QSFP, 100m with OM3 or 125m with OM4	QSFP-40G-SR-BD
Panduit Product Description	**Panduit Part Number**
QuickNet 24-Port Patch Panel	QPP24BL
QuickNet SFQ Series PanMPO to (6) LC Fiber Optic Cassette (OM4)	FQZO-12-10B1 &FQZO-12-10B2
QuickNet SFQ Series PanMPO to (6) LC Fiber Optic Cassette (OM3)	FQXO-12-10B1&FQXO-12-10B2
Opti-Core LC Interconnect Cable Assembly (OM4)	FZ2ERLNLNSN[**]
Opti-Core LC Interconnect Cable Assembly (OM3)	FX2ERLNLNSN[**]
QuickNet PanMPO Trunk Cable Assembly (OM4) (Configured with Female to Female, Method B[*])	FZTYP7E7EBAF[**]
QuickNet PanMPO Trunk Cable Assembly (OM3) (Configured with Female to Female, Method B[*])	FXTYP7E7EBAF[**]

[*] PanMPO cable assemblies can change polarity and gender as needed, but should be ordered with the polarity and gender configuration required to accelerate installation.

[**] Insert distance for correct Panduit product number (for example, a 20 foot cable is 020 orFZTYP7E7EBAF020).

Figure 27: Cisco Nexus 9000 Series BiDi Cabling Diagram
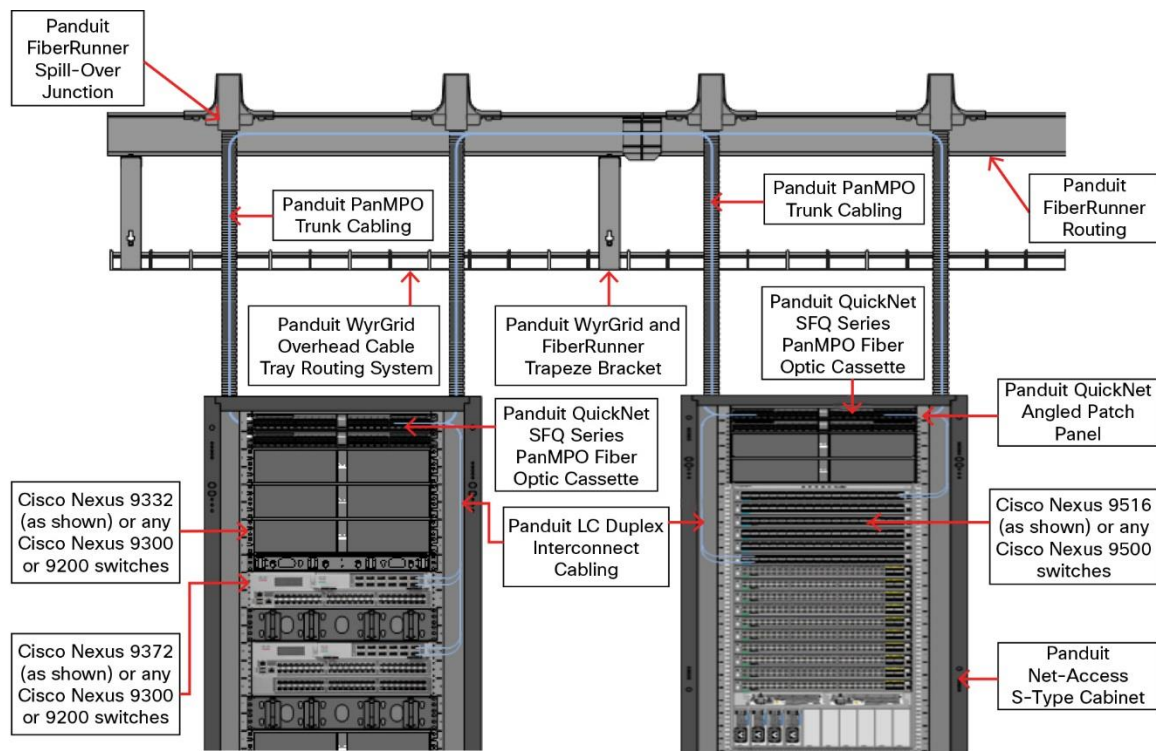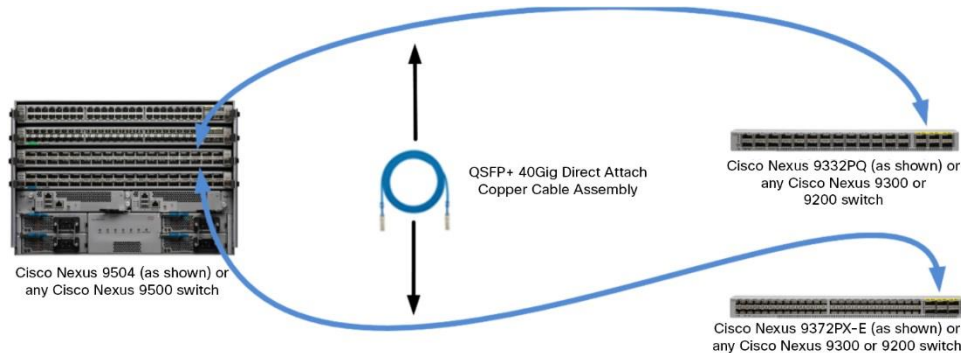
Copper Cabling Options
QSFP+ Direct-Attach Copper Cabling

Figure 28 shows a wiring option for a Cisco Nexus 9504 switch to a Cisco Nexus 9332PQ switch and a Cisco Nexus 9372PX-E switch using a QSFP+ direct-attach copper cable assembly. With a Panduit QSFP+ direct-attach copper cable assembly, this scenario is capable of distances of up to five meters. Table 7 shows the Cisco and Panduit Part Numbers for the direct-attach copper cable assemblies.

Figure 28: Cisco Nexus 9000 Series Switches QSFP+ 40Gbps Direct-Attach Copper Cabling Diagram

Table 7: Cisco Nexus 9000 Series Switches QSFP+ and SFP+ Direct-attach Copper Cabling Products

| Cisco Product Description	Cisco Part Number
Twinax Passive SFP+, 1m	SFP-H10GB-CU1M[1]
Twinax Passive SFP+, 1.5m	SFP-H10GB-CU1-5M[1]
Twinax Passive SFP+, 2m	SFP-H10GB-CU2M[1]
Twinax Passive SFP+, 2.5m	SFP-H10GB-CU2-5M[1]
Twinax Passive SFP+, 3m	SFP-H10GB-CU3M[1]
Twinax Passive SFP+, 5m	SFP-H10GB-CU5M[1]
40GBASE-CR4 direct-attach QSFP copper cable, 1 meter	QSFP-H40G-CU1M
40GBASE-CR4 direct-attach QSFP copper cable, 3 meter	QSFP-H40G-CU3M
40GBASE-CR4 direct-attach QSFP copper cable, 5 meter	QSFP-H40G-CU5M
Panduit Product Description	**Panduit Part Number**
SFP+ 10Gig Direct Attach Passive Cable Assembly, 1 to 7 Meters	PSF1PXA*MBU[1]
SFP+ 10Gig Direct Attach Active Cable Assembly, 7 to 15 Meters	PSF1AXD*MBU[1]
QSFP+ 40Gig Direct Attach Copper Cable Assembly, 1 Meter	PQSFPXA1MBU
QSFP+ 40Gig Direct Attach Copper Cable Assembly, 2 Meter	PQSFPXA2MBU
QSFP+ 40Gig Direct Attach Copper Cable Assembly, 3 Meter	PQSFPXA3MBU
QSFP+ 40Gig Direct Attach Copper Cable Assembly, 4 Meter	PQSFPXB4MBU
QSFP+ 40Gig Direct Attach Copper Cable Assembly, 5 Meter	PQSFPXC5MBU

[1] Will be supported in a future release.

* Insert distance for correct Panduit product number (for example, 1 meter cable is PSF1PXA1MBU).

QSFP+ to (4) SFP+ Hydra Cable Assembly

Figure 29 shows this scenario.

Figure 29: Cisco Nexus 9508 QSFP+ to SFP+ Hydra Cabling Diagram
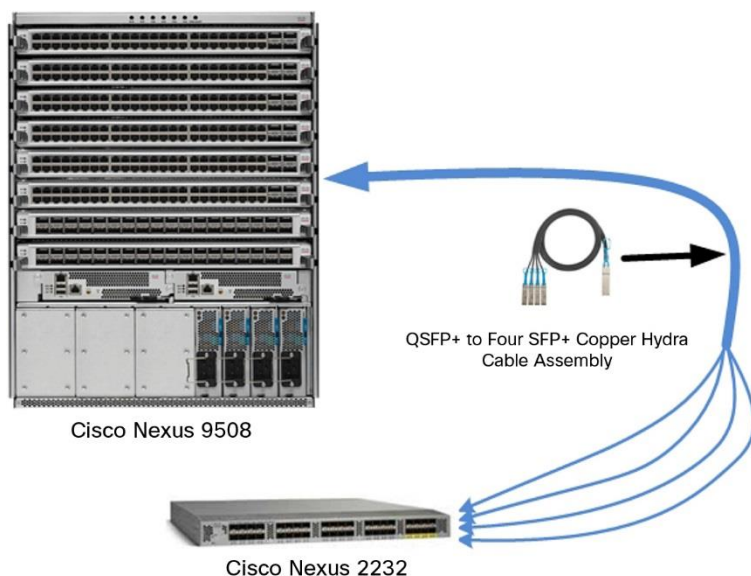
This scenario shows a wiring option for connecting a Cisco Nexus 9000 Series Switch to a Cisco Nexus 2232 using a QSFP+ to four SFP+ copper hydra cable assembly. By using passive copper cables, lengths up to five meters are possible. Active copper cables allow connectivity up to ten meters. This connection option is typical in ToR, MoR, and EoR deployments, where the distance between switches is within ten meters.

Table 8 lists the Cisco and Panduit cabling options with corresponding part numbers for QSFP+ to (4) SFP+ hydra cable assemblies.

Table 8: Cisco Nexus 9000 Series Switches QSFP+ to SFP+ Hydra Cabling Products

Cisco Product Description	Cisco Part Number
QSFP to 4x SFP+ direct-attach copper cable, 1 meter	QSFP-4SFP10G-CU1M
QSFP to 4x SFP+ direct-attach copper cable, 3 meter	QSFP-4SFP10G-CU3M
QSFP to 4x SFP+ direct-attach copper cable, 5 meter	QSFP-4SFP10G-CU5M
QSFP to 4 x SFP+ direct-attached Active Copper cable, 7 meters	QSFP-4x10G-AC7M
QSFP to 4 x SFP+ direct-attached Active Copper cable, 10 meters	QSFP-4x10G-AC10M
Panduit Product Description	Panduit Part Number
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 1 Meter	PHQ4SFPXA1MBL
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 1.5 Meter	PHQ4SFPXA1.5MBL
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 2 Meter	PHQ4SFPXA2MBL
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 2.5 Meter	PHQ4SFPXA2.5MBL
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 3 Meter	PHQ4SFPXA3MBL
QSFP+ to Four SFP+ Copper Hydra Cable Assembly, 3.5 Meter	PHQ4SFPXA3.5MBL

Panduit® Physical Infrastructure Solutions

Panduit offers a complete solution for the infrastructure supporting Cisco 10-40-100 Gbps switching platforms.

Thermal Management Solutions

In most data centers a hot aisle/cold aisle design is used to optimize the efficiency of the cooling system. Steps should be taken to preserve this separation, whenever possible. One such measure is thermal inlet ducting for Cisco Nexus 9300 and 9200 switches. The information below provides a guide for thermal inlet ducting products available from Panduit.

1RU Thermal Inlet Duct:

Figure 30 shows the thermal duct works with the following Cisco Nexus 9300 and 9200 switches:

9372
9332

Figure 30: Panduit 1RU Thermal Inlet Duct

2RU Thermal Inlet Duct:

Figure 31 shows the thermal duct works with the following Cisco Nexus 9300 and 9200 switches:

9396

Figure 31: Panduit 2RU Thermal Inlet Duct

3RU Thermal Inlet Duct:

Figure 32 shows the thermal duct works with the following Cisco Nexus 9300 and 9200 switches:

93128

Figure 32: Panduit 3RU Thermal Inlet Duct

Panduit® Cabinet and Cable Management Solutions

Panduit cabinet and cable management solutions help enable 10, 25, 40, and 100 Gigabit Ethernet switching platforms such as the Cisco Nexus 9000 Series Switches. Panduit cabinet and rack solutions support both two-post and four-post applications with sizes ranging from 42RU to 52RU. The cabinet solution also includes cable management and thermal management within and outside the cabinet. Panduit Thermal Ducting Solutions are specifically designed to work with the equipment used in a Cisco Nexus 9000 Series Switch deployment. Panduit Pre-configured Infrastructure Solutions simplify and accelerate deployment, enhance thermal performance, and decrease energy use. Figure 33 shows some of the features of Panduit cabinet and cable management component options that could make up its solutions.

Figure 33: Panduit Cabinet and Cable Management Components
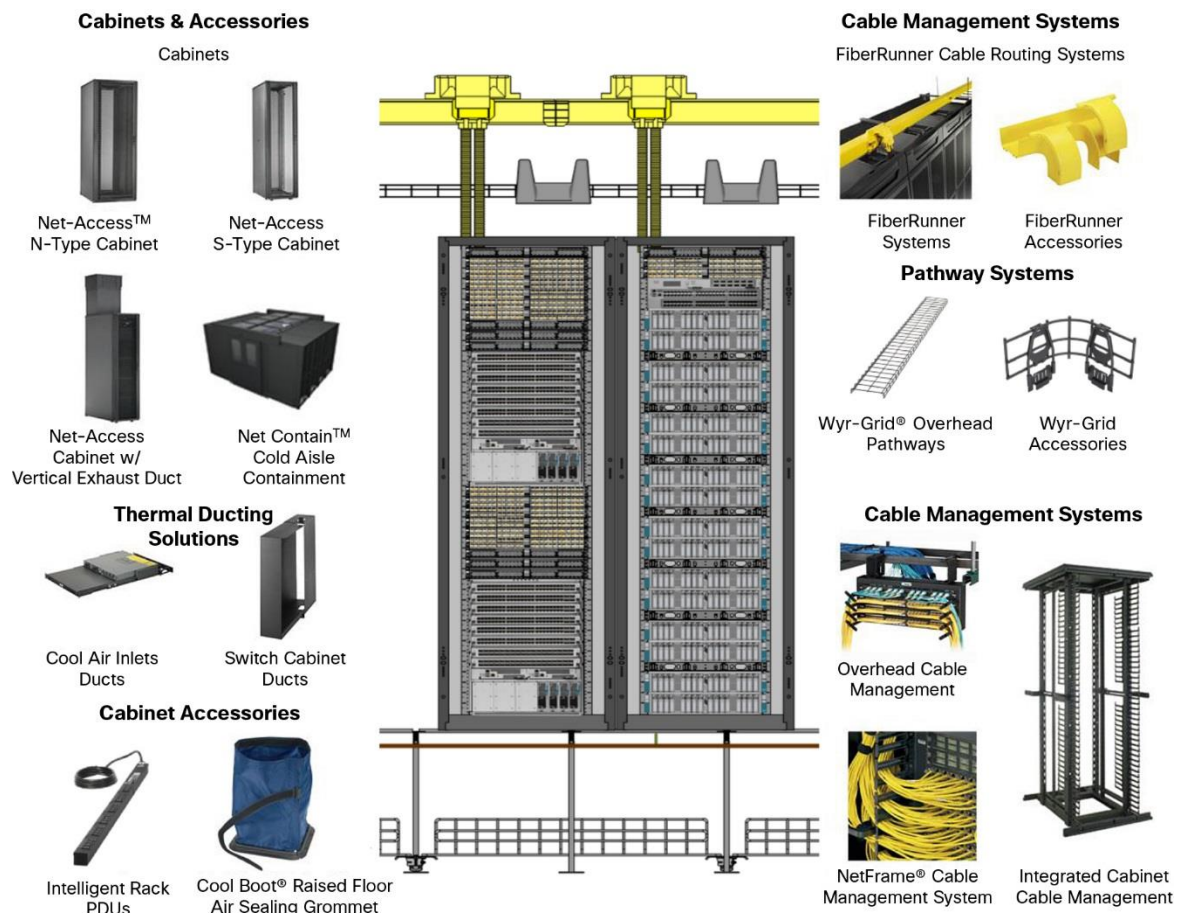
## Panduit® High-Speed Data Transport Solutions

Panduit fiber optic and copper cabling solutions for 10, 25, 40, and 100 Gigabit Ethernet support switching platforms such as the Cisco Nexus 9000 Series Switches. The Panduit Signature Core Fiber Optic Cabling System is a high-performance fiber that balances both modal and chromatic dispersion, allowing it to extend its reach beyond standard requirements. Panduit Signature Core extends the radius of fiber networks to meet customer implementation requirements. Panduit is leading the way in the relevant standard bodies to promote enhanced customer expectations through the use of this new technology.

Panduit copper technology incorporated into SFP+ 10-Gbps direct-attach copper (DAC) cable assemblies and TX6A 10Gig and TX6A-SD 10Gig UTP copper cable for 10GBASE-T provide enhanced performance characteristics. Both copper and fiber products can be incorporated into the Panduit QuickNet Cabling System, which offers factory-terminated and tested cable assemblies. This feature greatly accelerates the deployment of the solution and helps ensure expected performance.

Figure 34 shows some of the product options that can be used in a Panduit High-Speed Data Transport (HSDT) Solution for Cisco 9000 Series Switches.
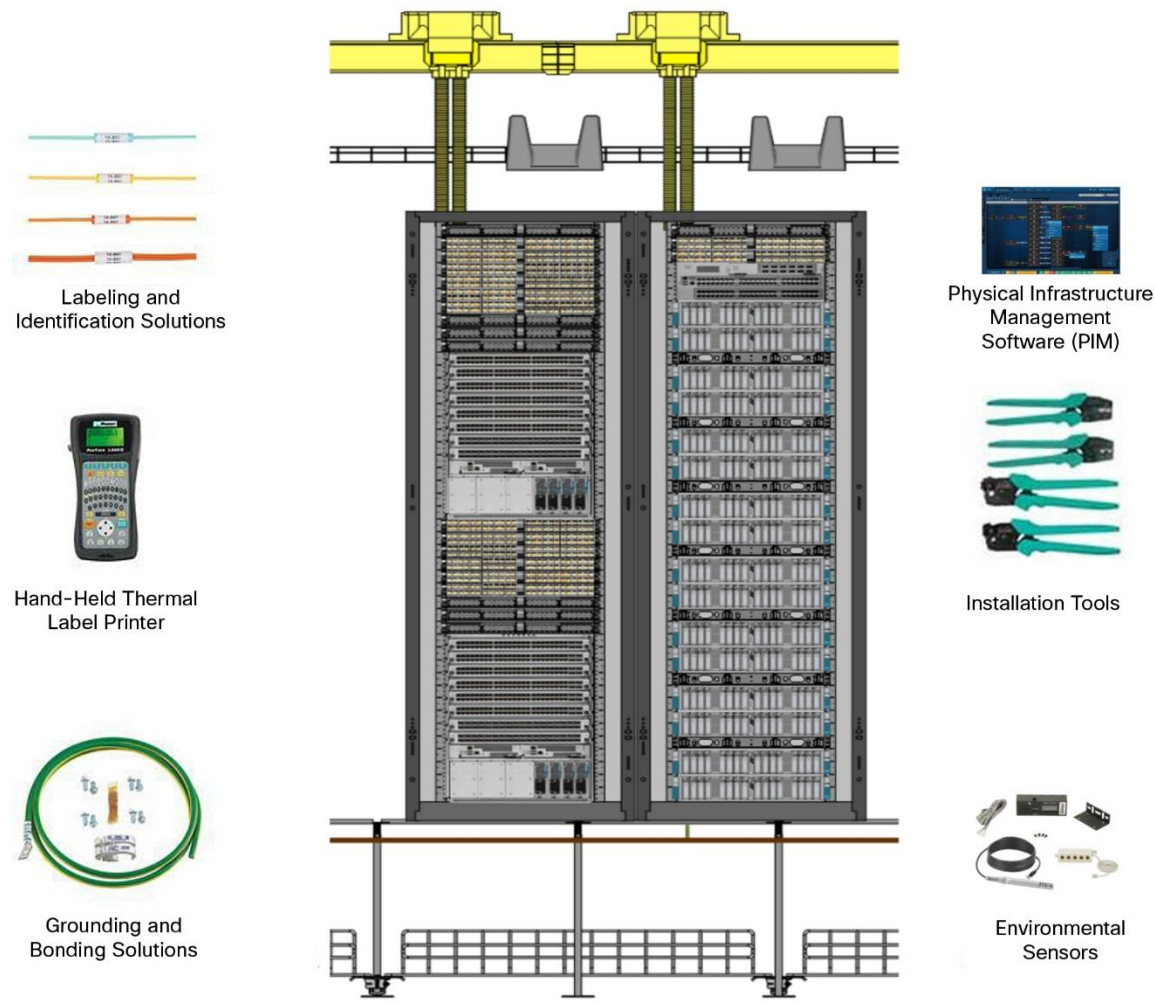
Figure 34: Panduit High-Speed Data Transport Components
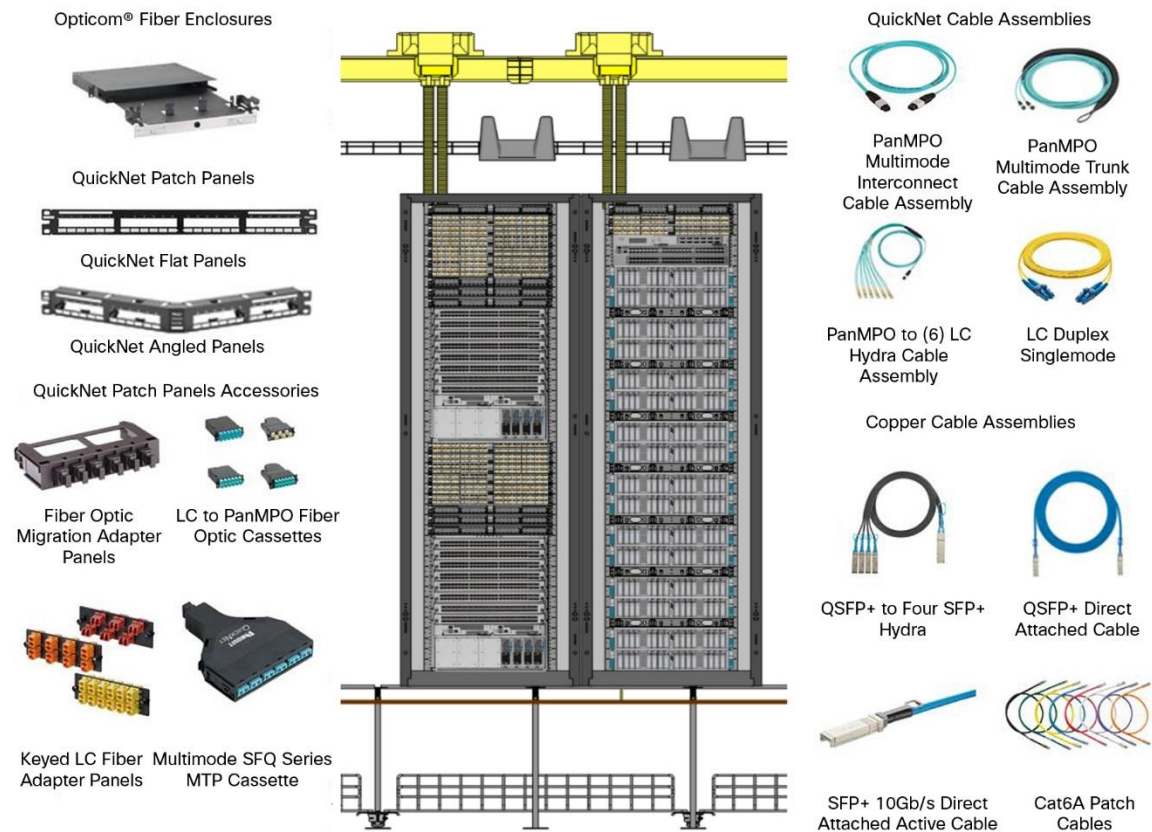
Panduit® Operation Management Solutions

Panduit operation management solutions help enable new technology such as the 10, 25, 40, and 100Gigabit Ethernet switching platform of the Cisco 9000 Series Switches. This includes features for grounding, identification and labeling of products, installation tools, and Panduit Physical Infrastructure Manager™ (PIM™) software and appliances. Panduit Advisory Services offer a broad array of assessment and optimization services to assist customers with the evaluation and deployment of Panduit Physical Infrastructure Solutions (Figure 35).

Figure 35: Panduit Operation Management Components

Labeling and
Identification Solutions

Hand-Held Thermal
Label Printer

Grounding and
Bonding Solutions

Physical Infrastructure
Management
Software (PIM)

Installation Tools

Environmental
Sensors

Conclusion

As virtualized environments are adopted by more enterprises, the pressure they apply on the network infrastructure will continue to cause data center stakeholders to reconsider the network architecture they use. The Cisco Nexus 9000 Series optimizes network performance to allow virtualized environments to function without the restrictions that traditional network architecture can impose. For a Cisco Nexus 9000 Series switching platform network to operate at the highest level, it requires a physical infrastructure built to help ensure superior performance. The Panduit Physical Infrastructure Solution gives Cisco the performance, flexibility, and reliability needed to operate without constraints.

For More Information

For more information, please contact Cisco at http://www.cisco.com or Panduit at http://www.panduit.com.

About Cisco

Cisco (NASDAQ: CSCO) is the worldwide leader in networking that transforms how people connect, communicate, and collaborate. Information about Cisco can be found at http://www.cisco.com. For ongoing news, please visit http://newsroom.cisco.com.

About Panduit

Panduit is a world leader that engineers flexible, end-to-end electrical and network connectivity physical infrastructure solutions that help business stay connected in a global world. Our high-performance products improve productivity and offer a lower total cost of ownership to create a competitive business advantage. Strong alliances with industry leaders, a global staff, and unmatched service and support, make Panduit a valuable, trusted partner. For more information, please visit http://www.panduit.com.

C11-729384-01 11/16