

Data Center Top-of-Rack Architecture Design

What You Will Learn

Forward-looking IT departments are preparing their data centers for the future by integrating support for 10 Gigabit Ethernet and a unified network fabric into their switching and cabling strategies. Since the typical data center lifecycle is 10 to 15 years, cabling architectures have a tremendous effect on the data center's ability to adapt to network architecture changes, evolving bandwidth needs, and technology moves, additions, and changes. Cabling architectures, if not chosen correctly, could force an early replacement of the cabling infrastructure to meet connectivity requirements as the network and computer technologies evolve.

Today's data centers deploy a variety of cabling models and architectures. With the migration from Gigabit Ethernet to 10 Gigabit Ethernet, cabling and network switching architectures are being reevaluated to help ensure a cost-effective and smooth data center transition. The choice of cabling architecture will affect throughput, expandability, sustainability, optimum density, energy management, total cost of ownership (TCO) and return on investment (ROI). Anticipating growth and technological changes can be difficult, but the data center should be able to respond to growth and changes in equipment, standards, and demands while remaining manageable and reliable.

This document examines the use of the top-of-rack (ToR) cabling and switching model for next-generation data center infrastructure. It explores current 10 Gigabit Ethernet cabling choices and provides a solution architecture based on ToR to address architectural challenges. Data center managers and facilities administrators will choose cabling architectures based on various factors. The ToR model offers a clear access-layer migration path to an optimized high-bandwidth network and cabling facilities architecture that features low capital and operating expenses and supports a "rack-and-roll" computer deployment model that increases business agility. The data center's access layer, or equipment distribution area (EDA), presents the biggest challenge to managers as they choose a cabling architecture to support data center computer connectivity needs. The ToR network architecture and cabling model proposes the use of fiber as the backbone cabling to the rack, with copper and fiber media for server connectivity at the rack level.

Introduction

The data center landscape is changing rapidly. IT departments building new data centers, expanding existing data center footprints, or updating racks of equipment all have to design a cabling and switching architecture that supports rapid change and mobility and accommodate transitions to 10, 40, and 100 Gigabit Ethernet over time. The main factors that IT departments must address include the following:

- **Modularity and flexibility is of paramount importance:** The need to rapidly deploy new applications and easily scale existing ones has caused server-at-a-time deployment to give way to a rack-at-a-time model. Many IT departments are ordering preconfigured racks of equipment with integrated cabling and switching and as many as 96 servers per rack. The time required to commission new racks and decommission old ones is now a matter of hours rather than days or weeks. Because different racks have different I/O requirements,

data center switching and cabling strategies must support a wide variety of connectivity requirements at any rack position.

- **Bandwidth requirements are increasing:** Today's powerful multsocket, multicore servers, blade systems, and integrated server and rack systems, often running virtualization software, are running at higher utilization levels and impose higher bandwidth demands. Some server racks are populated with servers requiring between five and seven Gigabit Ethernet connections and two Fibre Channel SAN connections each.
- **I/O connectivity options are evolving:** I/O connectivity options are evolving to accommodate the need for increasing bandwidth, and good data center switching and cabling strategies need to accommodate all connectivity requirements at any rack position. Racks today can be equipped with Gigabit Ethernet or 10 Gigabit Ethernet or a unified network fabric with Fibre Channel over Ethernet (FCoE).
- **Virtualization is being added at every layer of the data center:** Server virtualization is promoting server consolidation and increasing the need for bandwidth and access to network-attached storage (NAS). Virtualization is one of the main areas of focus for IT decision makers. Estimates suggest that the server virtualization market will grow by 44 percent over the next 4 years. The change to virtualization can be disruptive and necessitate a redesign of the networking infrastructure to gain all the benefits of the virtualized computer platform.

The challenge facing data centers today is how to support the modularity and flexibility that is needed to promote business agility and maintain a company's competitive edge. The same strategy that allows the intermixing of different rack types and I/O requirements must also support a varied set of connectivity options including Gigabit Ethernet and 10 Gigabit Ethernet as well as a unified network fabric.

Why Use Top-of-Rack Architecture

Rapidly changing business requirements impose a corresponding need for flexibility and mobility in data centers. Because of the significant cost of building a new data center, designing an infrastructure that provides the flexibility to meet business objectives while increasing ROI is an IT imperative. By building the data center infrastructure—power and cooling, cabling, etc.—in a modular fashion, data center flexibility can be increased, which in turn improves business agility.

Many organizations are now deploying modular data centers. IT departments are increasingly deploying not just servers but racks of servers at a time. Racks of servers, blade systems, and integrated rack-and-blade systems are often purchased in preconfigured racks with power, network, and storage cabling preinstalled so that racks can be commissioned within hours, not days, from the time they arrive on the loading dock. While server form factors are evolving, and some racks can host up to 96 independent computer resources, the rack form factor remains constant, making it the standard unit of deployment in many data centers.

ToR solutions complement rack-at-a-time deployment by simplifying and shortening cable runs and facilitating the replication of rack configurations. This rack-and-roll deployment model offers a solution by placing switching resources in each rack so that server connectivity can be aggregated and interconnected with the rest of the data center through a small number of cables connected to end-of-row (EoR) access- or aggregation-layer switches.

The TIA/EIA-942 specification provides a simple reference for data center cabling that supports different cabling schemes, EoR or ToR, to meet differing needs from a physical and operational perspective. The ToR model defines an architecture in which servers are connected to switches that are located within the same or adjacent racks, and in which these switches are connected to aggregation switches typically using horizontal fiber-optic cabling.

ToR switching allows oversubscription to be handled at the rack level, with a small number of fiber cables providing uniform connectivity to each rack. The advantage of this solution is that horizontal fiber can support different I/O connectivity options, including Gigabit Ethernet and 10 Gigabit Ethernet as well as Fibre Channel. The use of fiber from each rack also helps protect infrastructure investments as evolving standards, including 40 and 100 Gigabit Ethernet, are more likely to be implemented using fiber before any other transmission mechanism. By limiting the use of copper to within racks, the ToR model isolates the cabling that changes most often to the parts of the data center that change most frequently: the racks themselves. The use of fiber runs from racks provides a flexible data center cabling infrastructure that supports the transition from Gigabit Ethernet to 10 Gigabit Ethernet now, while enabling transition to 40 and 100 Gigabit Ethernet in the future.

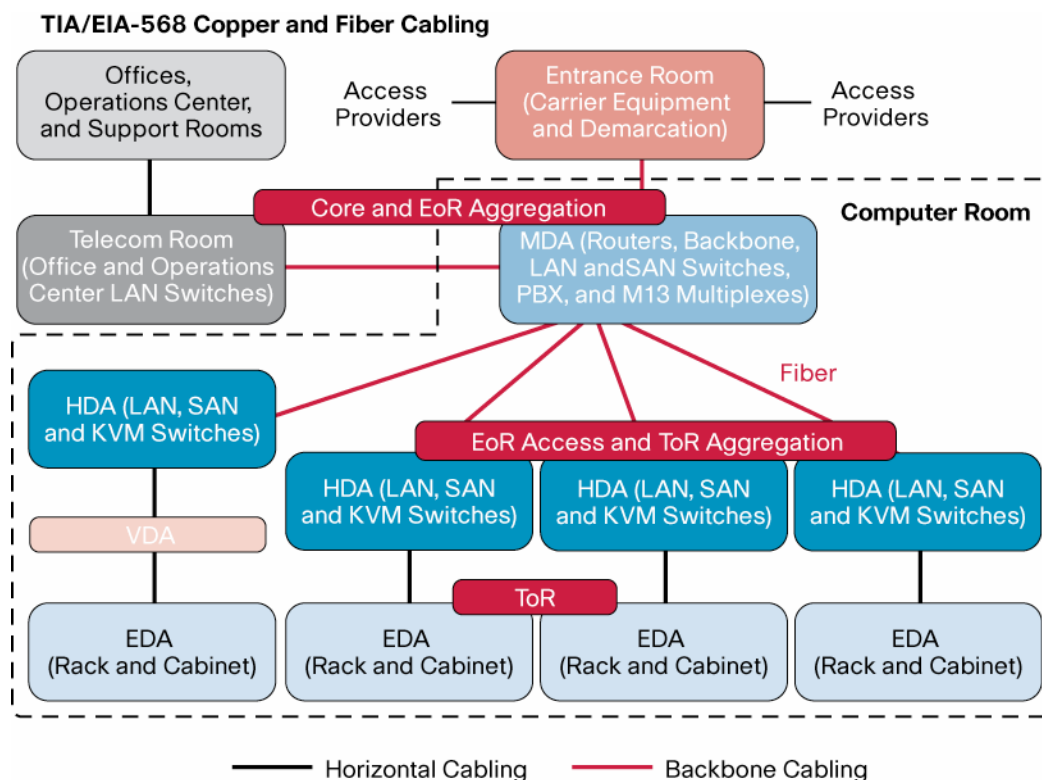
Modular Data Center Infrastructure

The cabling and infrastructure design for a modern data center is governed by multiple factors. The TIA/EIA-942 Telecommunications Infrastructure Standard for Data Centers provides guidelines for data center cabling infrastructure that customers can adopt as a guide in the data center cabling and planning process. Other standards, such as BICSI, provide guidelines for data center cabling and implementation. The TIA/EIA-942 cabling specification considers the need for flexibility, scalability, reliability, and space management (<http://www.tiaonline.org>). While the standard provides guidelines, specific design elements will vary with each data center. General considerations that apply to all data centers include:

- Support for storage devices (Fibre Channel, Small Computer System Interface [SCSI] or NAS, FCoE, etc.)
- Support for convergence and unified fabric with growth factors incorporated
- Reliability, scalability, and redundancy
- High-capacity and high-density server access requirements
- Flexibility and expandability with easy access for moves, additions, and changes
- Migration from Gigabit Ethernet to 10 Gigabit Ethernet server connectivity with future support for 40 and 100 Gigabit Ethernet
- Cabling architecture balance with power, cooling, structural loading, management, and operations needs

In the context of the TIA/EIA-942 simplified logical data center layout, the ToR architecture maps directly to the EDA and the horizontal distribution area (HDA). Figure 1 shows a mapping of the logical network architecture to the physical infrastructure.

Figure 1. TIA/EIA-942 Logical Layout



The EDA in the TIA/EIA-942 logical layout corresponds to the area where server racks are placed. Traditional structured copper cabling with a mix of fiber if needed for SAN or high-speed server connectivity is used to connect the EDA to the HDA. The environment requires careful planning to help ensure that the structured cabling meets the initial design requirements with enough room for growth. In cases where server racks are not yet in place or without the physical infrastructure needed to support rack-level flexibility, a zone distribution area (ZDA) cabling model is used. The ZDA allows for structured cabling to be placed under the floor or above the rack in anticipation of future server racks requiring connectivity to the network equipment that may be housed in the HDA. The ZDA follows a structured cabling model to the HDA. The primary difference in the horizontal cabling model between the ZDA and EDA is that the cables are terminated in the EDA racks, whereas the ZDA uses zone distribution blocks located outside the server racks.

The ToR cabling and network architecture in Figure 1 optimizes the requirement for horizontal cabling from the server rack by placing the ToR aggregation device at the top of the server rack. Actual placement of the ToR device may vary based on customer requirements (for example, in or above the server rack or ToR aggregation per two or three server racks) and seeks to optimize customer requirements for density, cabling, and design methodology. The aggregation of the ToR device will be in the HDA based on the specific TIA/EIA reference cabling model deployed. The ToR cabling design model follows a logical network layer construct in which the server network connectivity is aggregated at the rack level in the network access layer. The access layer is in turn connected to the network aggregation layer.

The ToR model uses fiber as the backbone cabling infrastructure that connects the EDA with the HDA and main distribution area (MDA). The ToR model augments the TIA/EIA-942 logical approach shown in Figure 1 by extending fiber as the backbone cabling of choice to the EDA or server rack. The amount of fiber required will vary based on design requirements. For example, in a migration-type design in which a ToR model is used for Ethernet aggregation and Fibre Channel connectivity is not unified I/O (UIO), additional fiber will be required for those servers requiring

Fibre Channel connectivity. In models in which UIO for Fibre Channel and Ethernet is assumed, the fiber requirement to the rack will be reduced. Other design augmentations, such as ToR for few racks with inter-rack cabling, will modify the fiber requirements. These are a few examples of some deployment models that have been adopted based on particular environment and requirements.

The Pod Approach

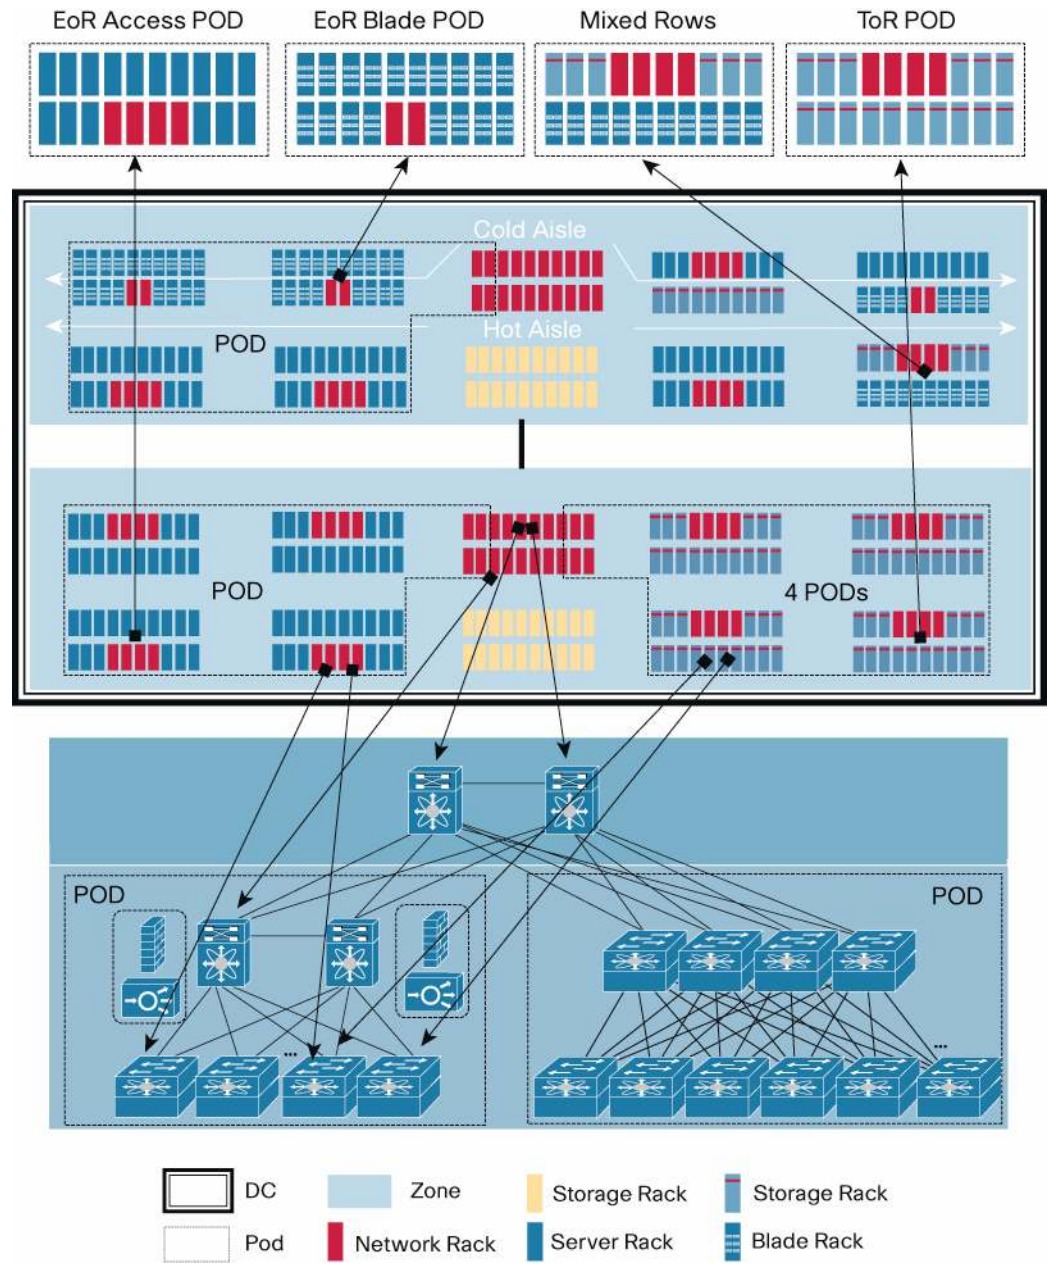
The easiest way to scale data center capacity is to use a highly modular architecture that enables the rapid deployment of infrastructure. One approach is to use a rack of servers as the base building block, with ToR switches, and with all copper cabling contained within the rack.

The TIA/EIA model provides a logical cabling layout that can be modularized for a data center implementation. When deploying large numbers of servers in the data center, it is extremely important that the design footprint be scalable. One way to simplify the design and simultaneously incorporate a scalable layout is to divide the data center floor space into modular, easily duplicated subareas.

This modular building block is used to design scalability into the network architecture at both OSI Layers 1 and 2. This design assumes that all computer resources incorporate resilient network, power, and storage resources, which translates to multiple LAN, SAN, and power connections within the physical layer infrastructure. The logical architecture is divided into three discrete layers (core, aggregation, and access), and the physical infrastructure is divided into manageable subareas called pods. From a network architecture perspective, a pod is defined by a pair of Ethernet aggregation switches. Pods support access-layer data link connectivity for low-latency interprocessor communications (IPC), unified fabric and unified IO, and Ethernet networks. A modular data center design coupled with a ToR access-layer architecture provides facilities and network architecture scalability, flexibility, and mobility to the rack level.

Figure 2 shows a modular layout for a Greenfield data center. The pod-based modular concept incorporates the various cabling architectures that may be required within the data center. More important, it enables a modular building block approach for server rack-and-roll deployment. In the reference architecture shown, the data center floor plan is divided into zones that are further subdivided into pods. A pod is used to represent a building block that defines the number of servers connected to a network aggregation block. Within the pod, server connectivity is handled at the rack level by the ToR device, which is in turn connected to the network aggregation layer.

Figure 2. Modular Data Center Layout



Top-of-Rack Access Media Considerations

When designing for 10-Gigabit Ethernet cabling infrastructure, the physics of the cabling plant, including signal attenuation, latency, and distance, as well as cabling installation and termination best practices, require careful consideration. Investment in the optimal cabling media for 10, 40, and 100 Gigabit Ethernet connectivity involves striking a balance among bandwidth, flexibility, and scalability. Table 1 shows the media and transceiver options for 10 Gigabit Ethernet connectivity available today.

Table 1. 10 Gigabit Ethernet Server Connectivity Options: UTP/F-UTP, MMF, SMF, Twinax, and CX4

Connector (Media)	Cable	Distance	Power (Each Side)	Transceiver Latency (Link)	Standard
SFP+ CU* copper	Twinax	<10m	~ 1.5W	~ .1 μ s	SFF 8431**
X2 CX4 copper	Twinax	<15m	4W	~ 0.1 μ s	IEEE 802.3ak
SFP+ USR MMF ultra-short reach	MM OM2 MM OM3	10m 100m	1W	~ 0	None
SFP+ SR MMF short reach	MM OM2 MM OM3	82m 300m	1W	~ 0	IEEE 802.3ae
RJ45 10GBASE-T copper	Cat6 Cat6a/7 Cat6a/7	55m 100m 30m	~ 6W*** ~ 6W*** ~ 4W***	2.5 μ s 2.5 μ s 1.5 μ s	IEEE 802.3an

* Terminated cable

** Draft 3.0, not final

*** As of 2008; expected to decrease over time

Several media and transceiver options are available today for 10 Gigabit Ethernet server connectivity. Considerations for use cases depend on variables such as cost, latency, distance, power consumption, and technology availability. For 10 Gigabit Ethernet server connectivity today, Small Form-Factor Pluggable (SFP) technology provides cost-effective and flexible options for server network interface cards (NICs) and ToR switches. SFP+ copper Twinax (SFP+cu) is currently the best choice for server connectivity.

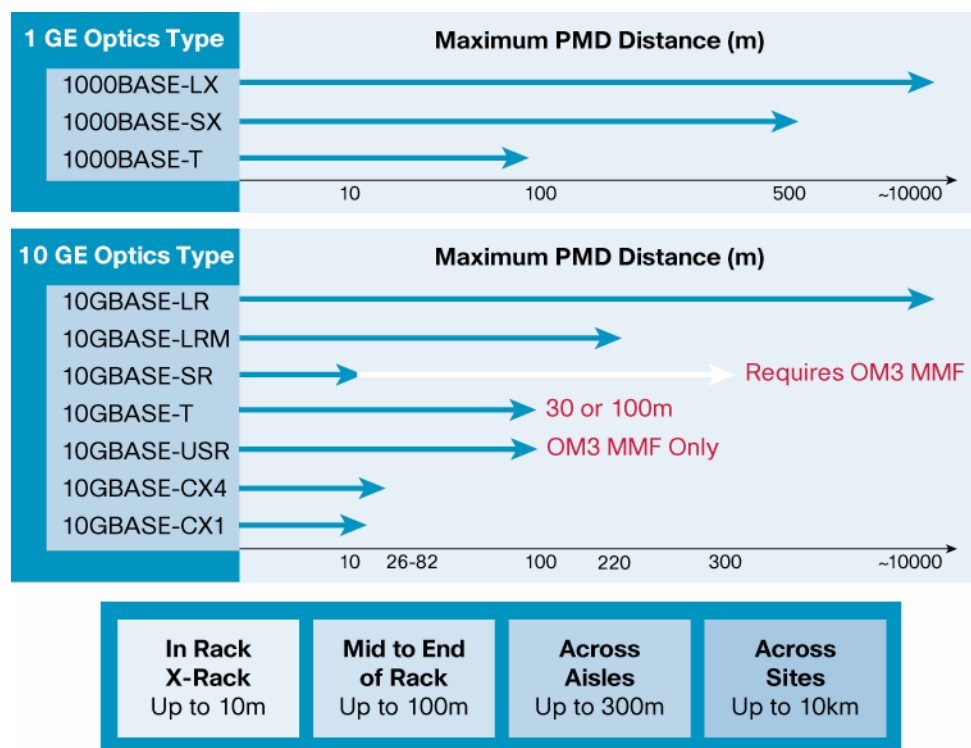
For EDA-to-HDA horizontal connectivity, SFP+ ultra-short-reach (USR) and short-reach (SR) cabling is better suited for longer distances between the server (EDA) rack and network equipment racks located in the HDA. SFP+ USR provides cost-effective fiber connectivity for fiber-based server connectivity options. SFP+ copper may also interconnect a number of EDA racks to a central EDA ToR device with horizontal backbone connectivity to the HDA or MDA.

10GBASE-T is currently power inefficient when compared to CX-1 and optical technologies. However, as more power-efficient third- and fourth-generation silicon architectures coupled with IEEE 802.3az Energy-Efficient Ethernet become available over the next few years, 10GBASE-T LAN on motherboard (LOM) and dense network switching products will become technically and economically viable server connectivity options.

Fiber and Copper Cabling Characteristics

A number of optics and media options are available for data center designs. Choices depend on the physical facilities requirements and media characteristics (physical medium dependent [PMD] factors) of the data center (Figure 3).

Figure 3. Media Options Based on Physical Requirements



Fiber

When considering a data center buildout, evaluation of the near- and mid-term future is required to maximize the cable plant investment. Because 10, 40, and 100 Gigabit Ethernet will be standardized in the next 3 to 5 years, organizations should assess fiber-optic cabling.

Several grades of high-bandwidth laser-optimized fiber are available for use in high-speed network installations, each with a different reach and data rate:

- 62.5/125-micrometer (OM1) fiber, designed to achieve 10- and 100-Mbps data rates; now largely an obsolete fiber.
- 50/125-micrometer (OM2) fiber, used to achieve 1-Gbps data rates
- 50/125-micrometer (OM2+, OM3, and OM3+) fiber, used to achieve 10-Gbps data rates; OM2+ and OM3+ fiber grades offer nearly double the bandwidth of their parent fibers (“+” indicates extended-reach OM2 and OM3 fiber)

Standard single-mode fiber (SMF) (ITU GR.652, TIA/EIA-493CAAA) is designed to support high-capacity, low-cost transmission components developed for the 1310-nanometer (nm) space. SMF features low dispersion and is optimized for use in the 1310-nm wavelength region. SMF is also used effectively with wavelength-division multiplexing (WDM) systems operating in the 15,550-nm wavelength region. SMF can be used for cross-aisle and inter-data center applications. SMF will support 40 and 100 Gigabit Ethernet with serial implementations.

The most cost-effective fiber cabling is multimode fiber (MMF). It is found in nearly all data centers today and can support 10 Gigabit Ethernet. However, as 40 and 100 Gigabit Ethernet become standardized, either MMF with parallel optics or SMF will be required. SMF, although simpler to manage and more amenable to higher bit-rate transmission than MMF, is more expensive to terminate and also more expensive in terms of the optical transceivers that are required at the switch and host devices.

Copper

Copper cabling can be used for 10 Gigabit Ethernet server connectivity.

Category 6A (Cat6a) copper cabling was developed in conjunction with the 10GBASE-T standard to achieve reliable 10-Gbps operation over 100-meter copper twisted-pair channels. Cat6a shielded and unshielded products are designed to extend usable bandwidth up to 500 MHz and to drastically reduce alien crosstalk interference. In the context of ToR architectures, Cat6a can be used within the rack.

1X-based Twinax copper* is an alternative copper solution that uses SFP+ direct-attach Twinax cabling. Although this solution has a limited cable distance of up to 10 meters, it provides a robust, power-efficient, and cost-effective solution for 10 Gigabit Ethernet transmission.

The SFP+ direct-attach solution is a fully integrated SFP+ cable that is available in multiple lengths up to 10 meters. As the cabling distance is limited, each server is directly connected to a ToR switch with no intermediate patch panels to manage. This approach dramatically simplifies cabling and termination as the cabling is contained within a single rack, and it works well with the concept of a modular data center. The SFP+ direct-attach solution draws 1.5 watts (W) of power per port, has a latency of 0.1 microsecond, and is available today.

Twinax cabling has been optimized for differential pair applications to support 10-Gbps signaling.¹

Top-of-Rack Cabling Recommendations

The ToR network architecture uses available cabling media options with flexibility at the rack level, using various server patch cable types, while taking advantage of fiber uplinks from the rack for horizontal cabling. Although CX-1 Twinax and fiber support high-speed 40- and 100-Gbps transmission, fiber is the recommended horizontal cabling media as it provides an optimal solution for high-speed 40- and 100-Gbps transmission over relatively long distances (up to 300 meters). Note that 40- and 100-Gbps transmission requires multiple fiber strands (OM3, OM4, and SMF fiber) plus optical interfaces that depend on the distance from the EDA to the HDA and aggregation.

Limiting the use of copper to within racks, the ToR model isolates the cabling that changes most frequently to the parts of the data center that change most frequently: the racks themselves. By using fiber runs from the racks, this architecture delivers a flexible cable infrastructure that supports the transition from Gigabit Ethernet to 10 Gigabit Ethernet today and allows adoption of 40 and 100 Gigabit Ethernet technologies in the future.


- In rack:
 - Cabling within the rack depends on connectivity requirements, which can be optimized for interface speed, latency, and cost of optics and transceiver
 - Copper options: SFP+ CX-1 Twinax (UTP, F/UTP, and S/FTP)
 - Fiber options: Low-cost SFP+ USR or SX, SR, and LM for short reach over MMF
- Outside rack (uplink to aggregation layer):

¹ It employs a unique homogeneous construction with 100 percent shielding that enables new levels of data rates over a single line with essentially no crosstalk. Both delay and amplitude skew are reduced because the one integral dielectric eliminates material variations and forces a nearly identical distance between conductors to be maintained. Twinax cabling is limited today to 10 meters for passive implementation and approximately 30 meters for active cabling solutions.

- Fiber (OM3 and OM4) if available; MMF recommended for lower cost on the fiber and the optics required for termination
- Fiber from rack (EDA) terminates at the aggregation layer (HDA)

Figure 4 shows a facilities rack view of a raised-floor server rack and network rack. The fiber-optic cabling is located over the server racks and terminates in fiber patch panels. Depending on requirements, the deployment may vary slightly, with the ToR device supporting a single rack or multiple racks.

Figure 4. Server and Network Racks



Top-of-Rack Architecture with the Cisco Nexus 5000 Series Switches and 2000 Series Fabric Extenders

The ToR architecture enables data center managers to implement a single cabling model that can support Gigabit Ethernet and 10 Gigabit Ethernet and unified network fabric today, while supporting future 40 and 100 Gigabit Ethernet standards as they come to market. Using a single overhead cable tray for fiber-optic cable management, data center managers have the flexibility to deploy preconfigured racks with different connectivity requirements in any rack position. For example, a rack of servers running multiple Gigabit Ethernet connections can be placed next to a rack of servers with 10 Gigabit Ethernet and FCoE connections to each server. As described here, Cisco Nexus™ products can be used to facilitate the ToR switching and cabling model.

Top-of-Rack Solution in Gigabit Ethernet Environments

Cisco® offers a compelling ToR solution that is supported by Cisco Nexus products. Using the Cisco Nexus 2148T Fabric Extender and Cisco Nexus 5000 Series Switches at the access layer, data centers can build self-contained racks of servers with Gigabit Ethernet connectivity using a small number of 10 Gigabit Ethernet fiber or CX-1 connections to an EoR or middle-of-row (MoR) switch.

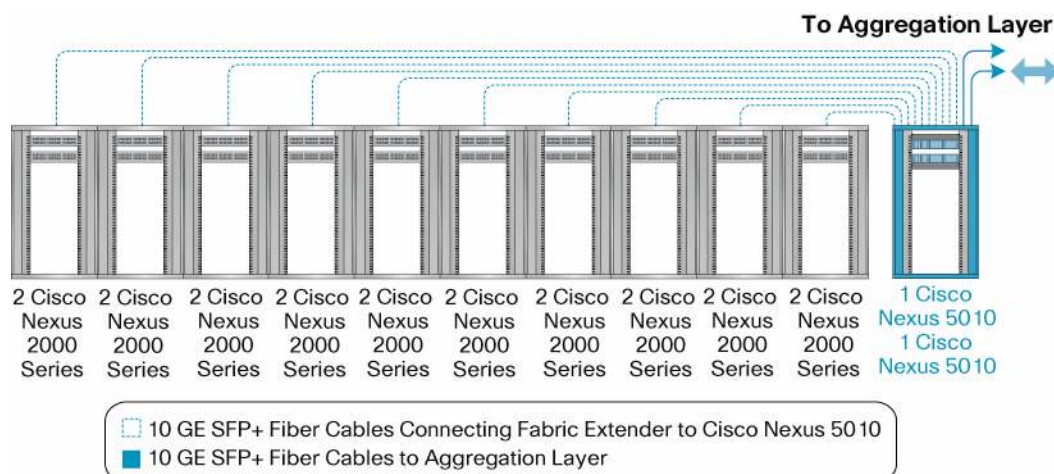
The Cisco Nexus 2148T is an innovative server aggregation mechanism that extends Cisco Nexus 5000 Series Switches into Gigabit Ethernet environments. Acting as a line card on a Cisco Nexus 5000 Series switch, the Cisco Nexus 2148T Fabric Extender aggregates up to 48 Gigabit Ethernet (fixed speed) connections at the rack and passes the network traffic up to the access-layer switch at the middle or end of the row. Because the Cisco Nexus fabric extender is an extension of the

switch itself, it offers massive scale with no increase in management complexity. Physically, it distributes the access layer across data center racks. Logically, the access layer remains at the end of the row and is managed by the Cisco Nexus 5000 Series switch. The Cisco Nexus fabric extender provides up to four SFP+ 10 Gigabit Ethernet uplinks supporting either 10 Gigabit Ethernet SFP+ optical transceivers or CX-1 direct-attach cable assemblies.

Figure 5 shows a row of dual-homed Gigabit Ethernet attached servers with two Cisco Nexus fabric extender devices in each rack that are connected to Cisco Nexus 5010 or 5020 Switches installed at the end of each row. Two Cisco Nexus fabric extenders in each rack support 40 servers per rack with no oversubscription if four 10 Gigabit Ethernet uplinks are configured.


Gigabit Ethernet attached servers are connected to the ToR Cisco Nexus fabric extender using Cat5e RJ-45 patch cables. Fiber uplinks provide connectivity between the fabric extender and upstream Cisco Nexus 5010 and 5020 Switches. If all four fabric extender uplinks are used to reduce oversubscription in the rack, then a total of 8 fiber strands are used from each rack to the end of the row.

Figure 5. Distributed Access Fabric with Nexus 5010 at End of Row



Short cable runs for 10 Gigabit Ethernet can be supported with SFP+ direct-attach copper (CX1) cables, which provide a low-cost, low-latency, and power-efficient server connection. SFP+ direct-attach 10 Gigabit Ethernet copper cables can connect the ToR Cisco Nexus fabric extender devices to the Cisco Nexus 5010 and 5020 Switches installed in the middle of the row. MoR configuration reduces the distance for the horizontal cable run, bringing it within the 7-meter range of CX1 cable. This layout can accommodate five or six 24-inch racks on either side of the Cisco Nexus 5020 and 5010 Switches installed in the middle of the row (Figure 6).

Figure 6. Distributed Access Fabric with Nexus 5010 at Middle of Row

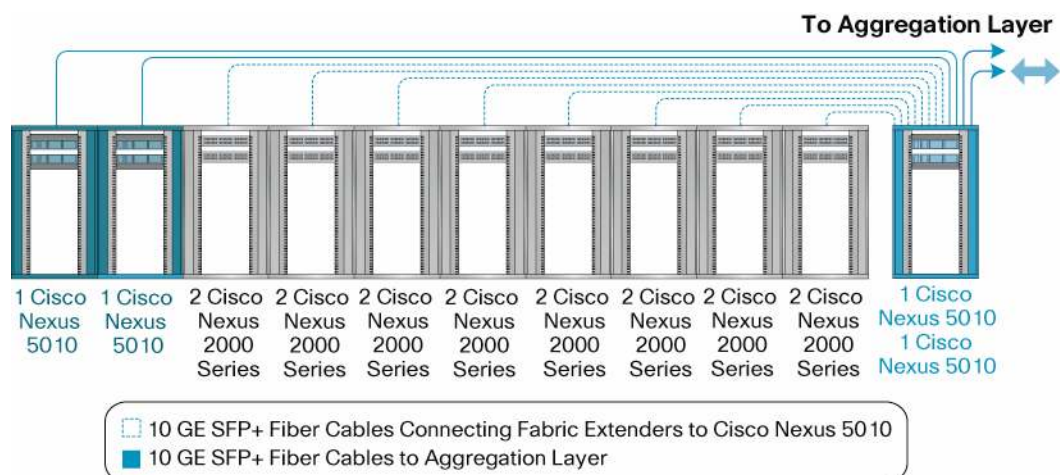


The Cisco Nexus 5010 is ideal for supporting 16 to 20 servers per rack in a ToR configuration. With 20 fixed SFP+ ports that support both fiber and SFP+ direct-attach 10 Gigabit Ethernet copper, each server can connect to the switch using CX-1 copper. If more than 20 ports are required, then a single expansion module can be installed to support up to 6 additional 10 Gigabit Ethernet connections. This architecture facilitates rack-and-roll deployment by allowing redundant networks to be constructed with two switches per rack and occupying a total of two rack units (2RU).

The Cisco Nexus 5000 Series Switches ease the transition to a unified network fabric by optionally carrying FCoE traffic from each server to the switch, and then native Fibre Channel from the switch to the SAN. Native Fibre Channel is supported through a choice of two expansion modules, one of which has eight Fibre Channel ports and one that has four 10 Gigabit Ethernet and four Fibre Channel ports.

Figure 7 shows that the same model of fiber in cable trays above racks and copper cabling within racks can support a mixed Gigabit Ethernet and 10 Gigabit Ethernet attached server environment within the same row. 10 Gigabit Ethernet attached servers are connected to the Cisco Nexus 5010 and 5020 ToR switches using direct-attach CX1 copper cables. Cisco Nexus 5010 and 5020 Switches are connected to the upstream aggregation-layer switch using fiber cables in the overhead cable tray. For the Gigabit Ethernet attached server racks, servers are connected to the Cisco Nexus ToR fabric extender using regular patch cords while fabric extenders are wired to the Cisco Nexus 5010 and 5020 EoR switches using fiber cables.

Figure 7. Distributed Access Fabric for 10GE and 1GE attached Servers



Conclusion

Data centers are undergoing a fundamental shift. Application developments and server virtualization are causing IT to seek a more flexible, efficient, and available infrastructure that can dynamically adapt to the needs of the business. To achieve this vision, data center architectures are changing to become more modular, so that data center managers can easily add racks of servers by adding racks that are precabled to a ToR switch and then connected to aggregation-layer switches using fiber-optic cables.

To deliver the necessary bandwidth, servers will transition to 10 Gigabit Ethernet in the short term for server and interswitch connectivity. The transition to 10 Gigabit Ethernet will also enable the adoption of a unified fabric using FCoE when organizations are ready to take advantage of the technology to achieve further efficiencies within the data center.

In the context of this data center evolution, cabling infrastructure must be considered to help ensure delivery of reliable high-bandwidth performance and application availability. Careful consideration also is needed to help increase power and cooling efficiency, as power and cooling have emerged as the top concerns for many organizations. Another consideration for Greenfield and Brownfield data center designs is the modular approach to data cabling and network architecture (this document refers to modules as pods). The modular approach improves data center design, deployment, and operations flexibility and increases efficiency and responsiveness, reducing overall capital expenditures (CapEx) and operating expenses (OpEx). The ToR architecture suits the modular approach and enables rack-and-roll deployment for further modularity and flexibility at the rack level. The ToR network and cabling architecture model enables a graceful transition from Gigabit Ethernet to 10 Gigabit Ethernet attached servers and provides infrastructure readiness for future adoption of 40 and 100 Gigabit Ethernet server and switch-to-switch connectivity without the need to change the cabling plant.

For More Information

<http://www.cisco.com/go/nexus5000>



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.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumina, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDE, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco ICS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FrameShare, Gigaset, HomeLink, Internet QuikStart, IOS, IPPhone, iQuik Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGC, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, SmartShare, SenderBase, SMI, SmartNet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet QuikStart, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (081215)