Introduction to Data Center

Top of Rack (ToR) Architecture Design

CISCO

Introduction to Data Center Top of Rack (ToR) Architecture Design

Table of Contents

Preface

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 life cycle is 10-15 years, structured cabling architectures have a tremendous impact on the data center's ability to adapt to network architecture changes, bandwidth and moves, adds, and changes (MACs). 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.

There are multiple cabling models and architectures deployed in today's data center. With the migration from 1GE to 10GE, cabling and network switching architectures are being reevaluated to ensure a cost-effective and smooth data center transition. The choice of cabling architecture will impact throughput, expandability, sustainability, optimum density, energy management, total cost of ownership (TCO) and return on investment (ROI). Anticipating growth and technological changes can be somewhat of a crystal ball prediction. The data center will take on a life of its own and should be able to respond to growth and changes in equipment, standards and demands all while remaining manageable and reliable.

This paper provides a look at the use of top-of-rack (ToR) cabling and switching model for next-generation data center infrastructure. The paper explores current 10G cabling choices and provides a solution architecture based on ToR to address architecture challenges. Data center managers and facilities folks 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 the data center computer connectivity needs. The ToR network architecture and cabling model, propose to use fiber as the backbone cabling to the rack with different 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-Gbps Ethernet over time. Key drivers of the changing landscape are:

- **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 multi-socket, multi-core 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 1 or 10 Gigabit Ethernet or a unified network fabric with Fibre Channel over Ethernet (FCoE).

- **Virtualization at every layer of the Data Center**. Server virtualization is driving server consolidation, greater bandwidth and access to network attached storage. Virtualization is one of the hottest areas for the IT decision makers. Estimates suggest that the server virtualization market will grow by 44% over the next four years. This change is often disruptive and necessitates a redesign of the networking infrastructure to realize 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 1- and 10-Gigabit Ethernet, as well as unified network fabric.

Top-of-the-Rack Architecture Drivers

Rapidly changing business requirements impose a corresponding need for flexibility and mobility in data centers. Due to the significant cost of building a new data center, designing an infrastructure that provides the flexibility to meet business objectives while maximizing return on investment 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 *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 compliment rack-at-a-time deployment by simplifying and shortening cable runs and easing the replication of rack configurations. This "rack-and-roll" deployment model offers a solution by placing switching resources into 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 access or aggregation-layer switches.

The TIA 942 specification provides a simple reference for data center cabling that supports different cabling schemes – end-of-row or top-of-rack – to meet differing needs from a physical and operational perspective. One model, top-of-rack defines an architecture where servers are connected to switches that are located within the same or adjacent racks, and where these switches are connected to aggregation switches typically using horizontal fiber optic cabling.

Top-of-rack switching allows oversubscription to be handled at the rack level, with a small number of fiber cables providing uniform connectivity to each rack. The beauty of this solution is that horizontal fiber can support different I/O connectivity options, including 1- and 10-Gigabit Ethernet as well as Fibre Channel. The use of fiber from each rack is also a form of future-proofing as evolving standards, including 40- and 100-Gbps 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. By using fiber runs from racks, a flexible data center cabling infrastructure that supports the transition

from 1- to 10-Gigabit Ethernet now, while providing a future-proof solution that enables transition to 40- and 100-Gbps Ethernet in the future.

Modular Data Center Infrastructure

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

- Support for storage devices (i.e. Fibre channel, SCSI or NAS, FCoE)

- Support for convergence/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, adds and changes

- Migration consideration for 1GE to 10GE server connectivity with future-proofing to support 40G and 100G.

- Cabling architecture balance with power, cooling, structural loading, management and operations

In the context of the TIA-942 data center simplified logical layout, the ToR architecture will map directly to the Equipment Distribution Area (EDA) and the Horizontal Distribution Area (HDA). The diagram below illustrates a mapping of the logical network architecture to the physical infrastructure.

TIA-942 Logical Layout
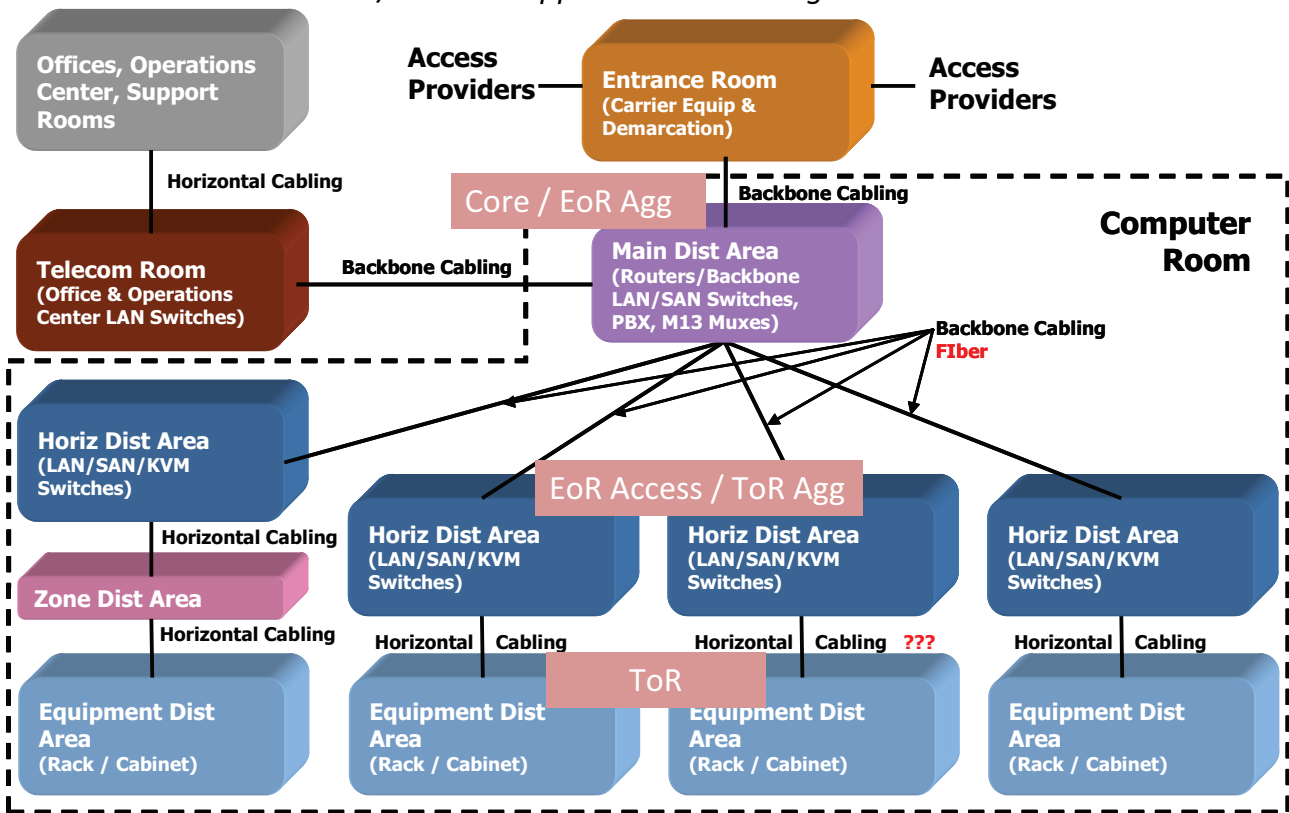
EIA/TIA 568 Copper & Fiber Cabling

Fig. 1.1

The equipment distribution area (EDA) in the TIA-942 logical layout would correspond 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 would be used to connect the EDA to the HDA. The environment would require careful planning to ensure that the structured cabling will meet the initial design requirements with enough room for growth. In the cases where server racks were not yet in place or the physical infrastructure needed to support rack level flexibility, a Zone Distribution Area cabling model would be used. The zone distribution area allows for structured cabling to be placed under floor or above 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 of horizontal cabling model between the ZDA and EDA is that the cables are terminated in the EDA racks were as the ZDA uses zone distribution blocks located outside the server racks.

The ToR cabling and network architecture in the context figure 1.1 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, (e.g. either 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/switch will be in the HDA based on the specific TIA reference cabling model deployed The ToR cabling design model follows a logical network layer construct where 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.

Connectivity from EDA to HDA the ToR model will use fiber as the backbone cabling infrastructure that connects the EDA with the HAD and MDA. The ToR model augments the TIA/EIA-942 logical approach shown above in Fig. 1.1 by extending fiber as the backbone cabling of choice to the EDA/server rack. The amount of fiber required will vary based on design requirements. For example, in a migration-type design where a ToR model is used for Ethernet aggregation and Fibre Channel (FC) connectivity is non-unified I/O (UIO), then additional fiber will be required for those servers requiring FC connectivity. In models where UIO for FC and Ethernet is assumed, the fiber requirement to the rack will be reduced. Other design augmentation, like a 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 the 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, where all copper cabling is contained within the rack.

The TIA model provides a logical cabling layout that can be modularized for a data center build out. When deploying large volumes of servers inside 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 sub-areas.

This modular building block is used to design scalability into the network architecture at both OSI Layers 1 and 2. It is assumed that all computer resources incorporate resilient network, power, and storage resources. This assumption translates to multiple LAN, SAN, and power connections within the physical layer infrastructure. The logical architecture is divided into three discrete layers (Core, Aggregation, Access), and the physical infrastructure is divided into manageable sub-areas called *pod*s. 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 inter-processor communications (IPC), Unified Fabric/Unified IO and Ethernet networks. A modular data center build coupled with a ToR access layer architecture provides facilities and network architecture scalability, flexibility and mobility to the rack level.

Physical Infrastructure and Network Topology
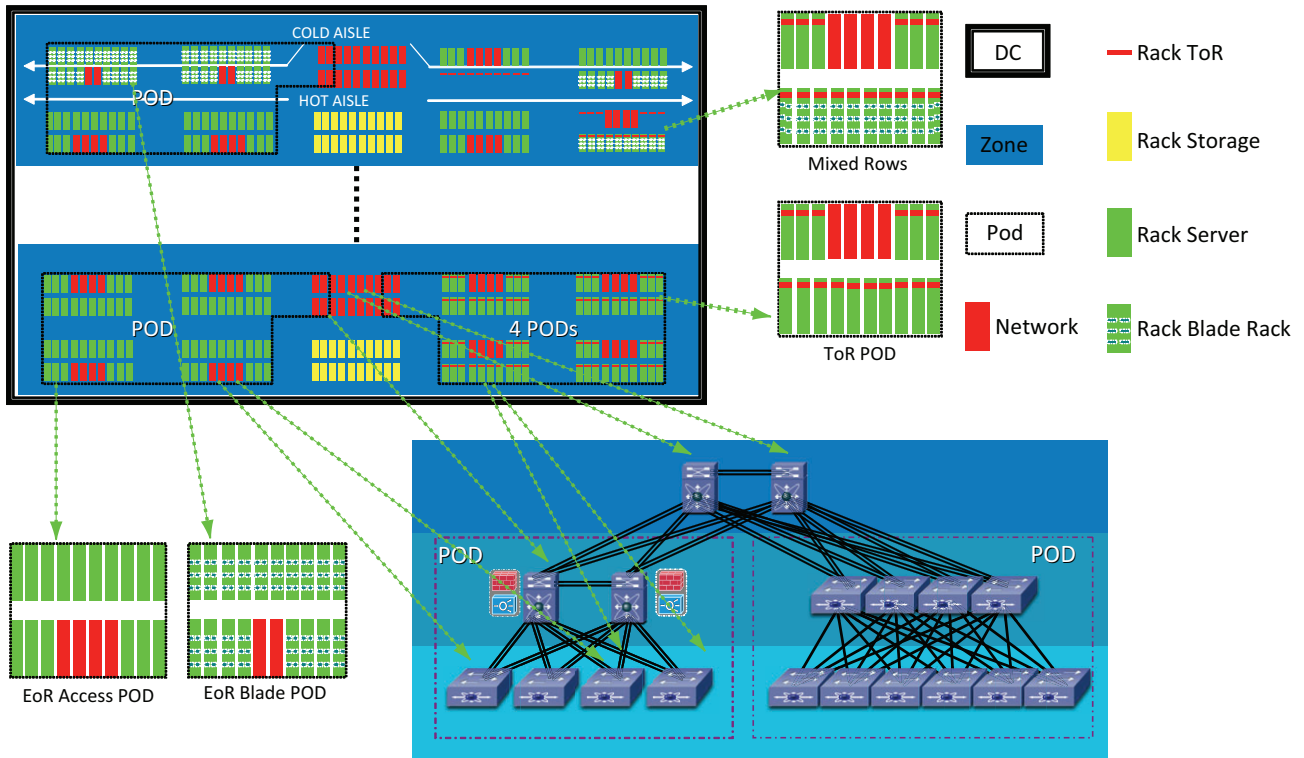## *Modular Cabling Architecture Methodology*

Fig. 1.2

Figure 1.2 above shows a modular layout for a Greenfield data center. The pod-based modular concept incorporates the different cabling architectures that may be required within the data center. More importantly it enables a modular building block approach for server "rack and roll" deployment. In the reference 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.

Top of Rack Access Media Considerations

When considering 10-Gigabit Ethernet cabling infrastructure, the physics of the cabling plant must be considered. Factors such as signal attenuation, latency and distance, as well as cabling installation and termination best practices, are factors for careful consideration. Investing in the optimal cabling media for 10-, 40- and 100-Gigabit connectivity involves striking a balance between bandwidth, flexibility and scalability. Figure 1.3 below shows media and transceiver options for 10G connectivity available today.

10G Sever Connectivity Options–UTP/F-UTP, MMF, SMF, TwinAx, CX4

Connector (Media)	Cable	Distance	Power (each side)	Transceiver Latency (link)	Standard
SFP+ CU* copper	Twinax	<10m	~ 1.5W	~ .1 μs	SFF 8431**
X2 CX4 copper	Twinax	15m	4W	~ 0.1 μs	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 μs 2.5 μs 1.5 μs	IEEE 802.3an

* Terminated cable   ** Draft 3.0, not final   *** As of 2008; expected to decrease over time

Fig. 1.3

There are several media and transceiver options available today for 10G server connectivity. Considerations for use case depend on variables like cost, latency, distance, power consumption and technology availability. For 10G server connectivity today, SFP-based technology provides cost-effective and flexible options for server NIC and ToR switches. SFP+cu (copper twinax) is currently the most optimal choice for server connectivity.

For EDA to HDA horizontal connectivity, SFP+ USR/SR is better suited for longer distances between 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.

10GBaseT is, as of time of writing, power inefficient when compared to CX-1 and optical technologies. However, as more power efficient 3rd- and 4th-generation silicon architectures coupled with IEEE 802.3az Energy Efficient Ethernet become available over the next few years, 10GBASE-T LOM (LAN on motherboard) and dense network switching products will become technically and economically viable sever connectivity option.

Fiber and Copper Cabling Characteristics

For data center builds there are a number of choices in optics and media. Consideration is required relative to the data center physical facilities requirements and media characteristics. The table below provides guidance on different optics and usage in the data center.
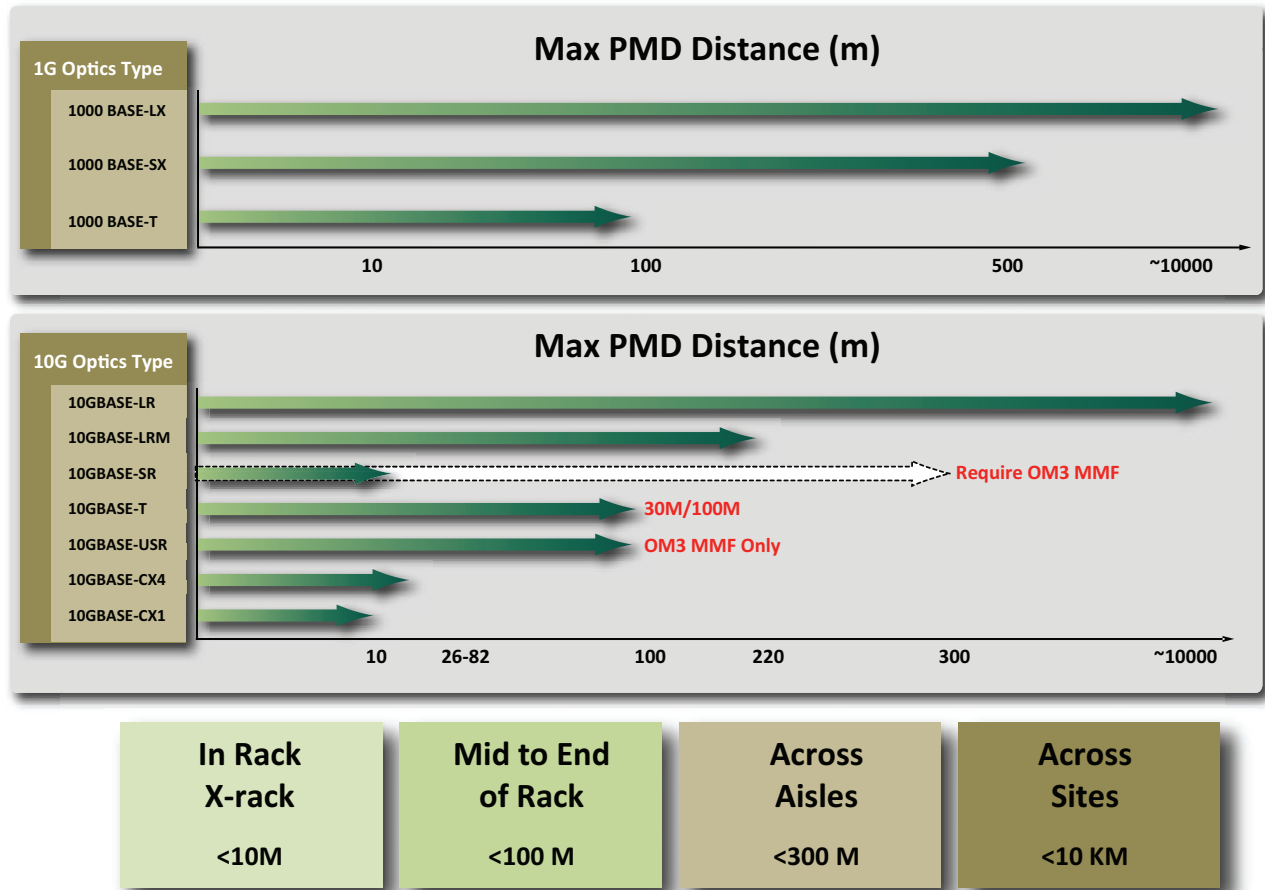
Max PMD Distance (m)

| 1G Optics Type |
| 1000 BASE-LX |
| 1000 BASE-SX |
| 1000 BASE-T |

10   100   500   ~10000

Max PMD Distance (m)

| 10G Optics Type |
| 10GBASE-LR |
| 10GBASE-LRM |
| 10GBASE-SR — Require OM3 MMF |
| 10GBASE-T — 30M/100M |
| 10GBASE-USR — OM3 MMF Only |
| 10GBASE-CX4 |
| 10GBASE-CX1 |

10   26-82   100   220   300   ~10000

In Rack X-rack	Mid to End of Rack	Across Aisles	Across Sites
<10M	<100 M	<300 M	<10 KM

Fig. 1.4

Fiber

When considering a data center build out, evaluation of near- and mid-term future is required to maximize the investment in cable plant. Because 10-, 40- and 100-Gigabit Ethernet will be standardized in the next 3-5 years, one 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:

- **62.5/125μm (OM1) fiber**, designed to achieve 10 Mb/s and 100 Mb/s data rates, is now largely a legacy fiber

- **50/125μm (OM2) fiber**, used to achieve 1 Gb/s data rates

- **50/125μm (OM2+, OM3, and OM3+) fiber**, *used to achieve 10 Gb/s data rates; OM2+ and OM3+ fiber grades* offer nearly double the bandwidth of their parent fibers ('+' represents extended reach OM2 and OM3 fiber).

- **SMF (ITU GR.652, TIA/EIA-493CAAA),** standard single-mode fiber. Designed to support high capacity, low-cost transmission components developed for the 1310 nm window. SMF fibers feature low dispersion and is optimized for use in the 1310 nm wavelength region. SMF fiber is also used effectively with WDM systems operating in the 15550 nm wavelength region.. SMF fiber can be used for cross-isle and inter-data center applications. SMF fiber will support 40, 100G with serial implementations. (update this section)

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

Copper

Copper cabling consideration for 10G server connectivity.

- **Category 6A copper** cabling was developed in conjunction with the 10GBASE-T standard to achieve reliable 10 Gb/s operation over 100 m copper twisted-pair channels. Category 6A 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.l (describe 6a copper more)

- **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+ and cable that is available in multiple lengths up to 10 meters. As the cabling distance is limited, each server is directly connected to a top-of-rack switch with no intermediate patch panels to manage. This dramatically simplifies cabling and termination as the cabling is contained within a single rack and works well with the concept of a modular data center. The SFP+ directly attached solution draws 0.1 Watts power per port, has a latency of 0.1 microseconds, and is available today.

* Twin-ax cabling has been optimized for differential pair applications to support 10 Gb/s signaling. It employs a unique homogeneous construction with 100% shielding that enables completely new levels of data rates over a single line with virtually no crosstalk. Both delay and amplitude skew are minimized 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 ~30 meters for active cabling solutions.

Top-of-the-Rack Cabling Recommendations

The ToR network architecture leverages available cabling media options with flexibility at the rack level to utilize different sever 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-G transmission, fiber is the recommended horizontal cabling media as it provides an optimal solution for high-speed 40- and 100-G transmission over relatively long distances (up to 300 meters). It should be noted that 40- and 100-G transmission requires multiple fiber strands (OM3/OM4/SMF fiber) plus optical interfaces that depend on the distance from EDA to HDA / aggregation.

Limiting the use of copper to within racks, the ToR model isolates the cabling that changes most rapidly 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 future-proof cable infrastructure that supports the transition from 1-Gigabit Ethernet to 10-Gigabit Ethernet today, and allows adoption of 40- and 100-Gbps Ethernet technologies in the future.

InRack: Cabling within the rack is dependent on connectivity requirements that may be optimized around
- interface speed, latency, and cost of optics/transceiver.

- Copper options: SFP+ CX-1 Twinax, (UTP, F/UTP, S/FTP) or

- Fiber options: low-cost SFP+ USR or SX, SR, LM for short reach over multimode fiber

Outside Rack: (Uplink to Aggregation)
- Fiber (OM3/OM4) if available; MM fiber recommended for lower cost on the fiber and the optics required for termination.

- Fiber from rack (EDA) terminates at the aggregation (HDA)

Shown: OM3 MM Fiber per Rack, terminating on patch panels in Network Rack
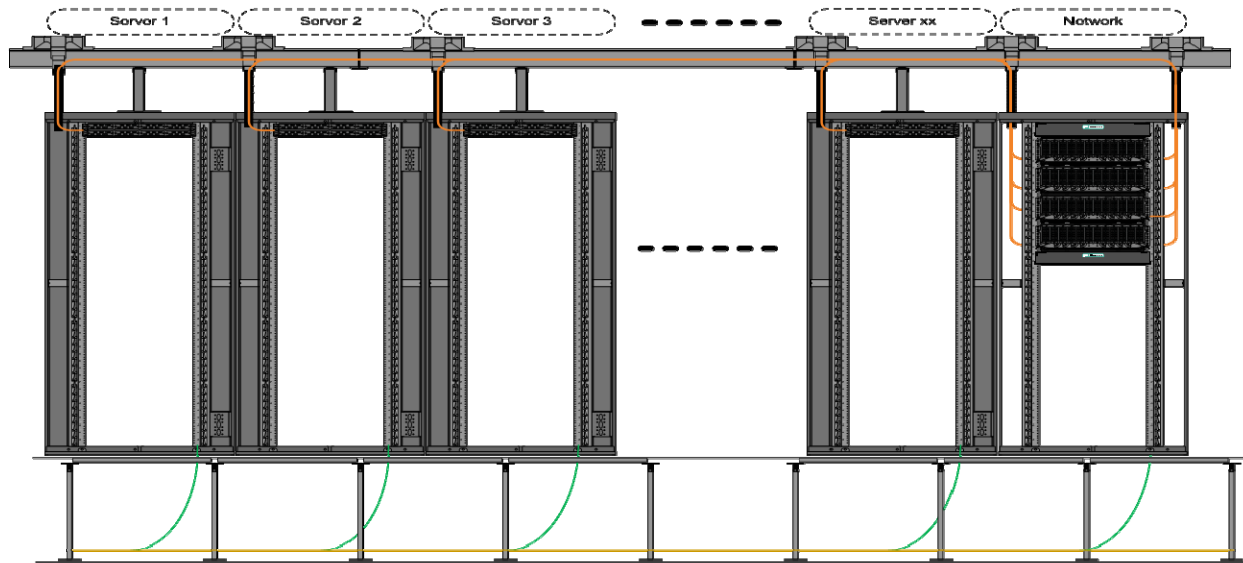*Fiber count vary based on design requirement (SAN, ToR, Blade etc..)*

Fig. 1.5

Figure 1.5 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 in terms of the ToR device supporting a single rack or multiple racks.

Top-of-the-Rack Model at Work with the Nexus 5000 and 2000

The ToR architecture offers data center managers the ability to implement a single cabling model that can support 1- and 10-Gigabit Ethernet and unified network fabric today, while supporting future 40- and 100-Gbps 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 might be placed next to a rack of servers with 10-Gigabit Ethernet and Fibre Channel over Ethernet connections to each server. This section of the paper demonstrates how Cisco Nexus products facilitate the top-of-rack switching and cabling model.

Top of the Rack in Gigabit Ethernet Environments

Cisco offers a compelling ToR solution that is supported by Cisco Nexus™ products. Using the Cisco Nexus 2148-T and Nexus 5000 switches at the access layer, data centers can build self-contained racks of servers with Gigabit Ethernet

connectivity requirements using a small number of 10-Gigabit Ethernet fiber or CX-1 connections to an end-of-row or middle-of-row switch.

The Cisco Nexus 2148-T (Fabric Extender) 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 Fabric Extender aggregates up to forty-eight 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+ 10G uplinks supporting either 10G SFP+ optical transceivers or CX-1 direct-attach cable assemblies.

Figure 1.6 below illustrates a row of dual-homed 1-Gig attached servers with two Nexus Fabric Extender devices in each rack that are connected to Nexus 5010 or 5020 switches installed at the end of each row. Two Cisco Nexus Fabric Extenders in each rack support forty servers per rack with no oversubscription if four 10-Gigabit Ethernet uplinks are configured.

1-G attached servers are connected to the ToR Nexus Fabric Extender using Cat5e RJ-45 patch cables. Fibre uplinks provide connectivity between Nexus Fabric Extender and upstream Nexus 5010/ 5020 switches. If all four Fabric Extender uplinks are used to minimize oversubscription in the rack, then a total of 8 Fibre strands are utilized from each rack to the end of the row.
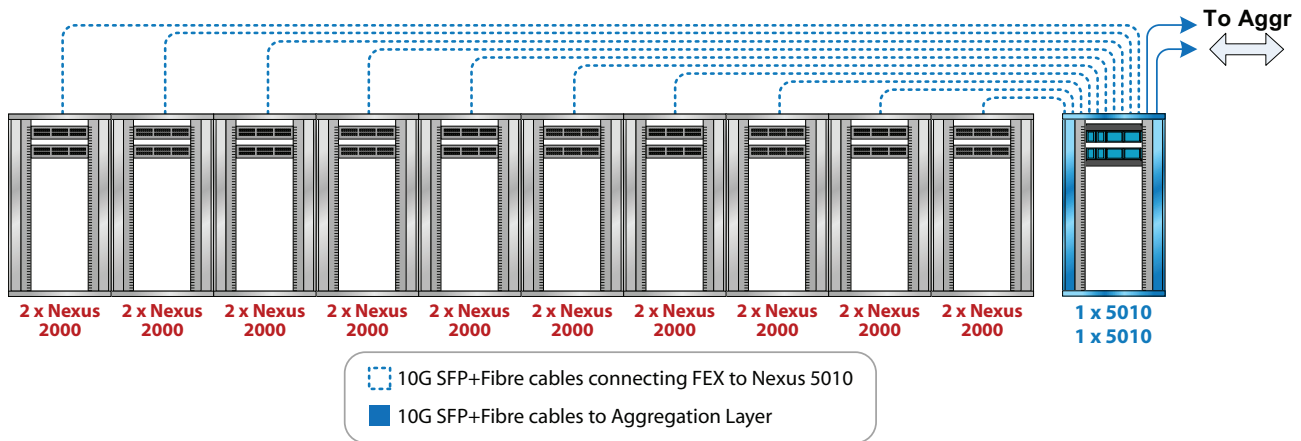

Fig. 1.6

Short cable runs for 10-Gigabit Ethernet can be supported with SFP+ direct-attached copper (CX1) cables that provide a low-cost, low-latency and power-efficient server connection. SFP+ direct-attached 10G copper cables can connect the ToR Nexus Fabric Extender devices to the Nexus 5010/ 5020 switches installed in the middle of row. Middle of row configuration reduces the distance for the horizontal cable run, bringing it within the 7m range of CX1 cable. This layout can accommodate five to six 24-inch racks on either side of the Nexus 5020/5010 switches installed in the middle of the row.

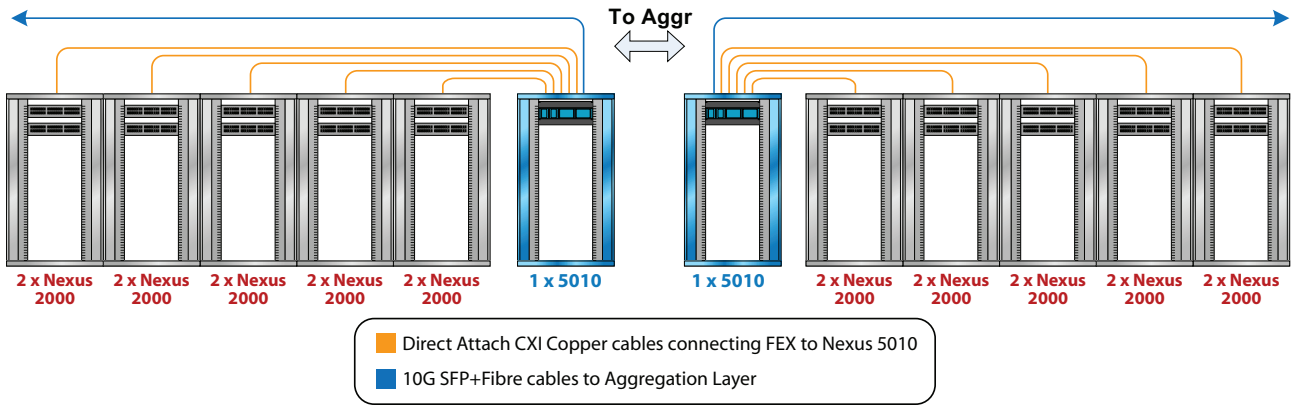[*] *Nexus 2148-T supports only fixed-speed 1-G interfaces*

Fig. 1.7

The Cisco Nexus 5010 Switch is ideal for supporting sixteen to twenty servers per rack in a ToR configuration. With twenty fixed SFP+ ports that support both Fibre and SFP+ direct-attached 10-Gigabit copper, each server can connect to the switch using CX-1 copper. If more than twenty ports are required, then a single expansion module can be installed to support up to six 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 2 RU.

The Cisco Nexus 5000 Series Switch eases 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.

The following figure illustrates that the same model of fiber in cable trays above racks and copper cabling within racks can support a mixed 1-G and 10-G attached server environment within the same row. 10-G attached servers are connected to the ToR Nexus 5010/5020 switches using direct-attach CX1 copper cables. Nexus 5010/5020 switches are connected to the upstream aggregating layer switch using Fibre cables in the overhead cable tray. For the 1-G attached-server racks, servers are connected to the ToR Nexus Fabric Extender using regular patch cords while Fabric Extenders are wired to the end-of-row Nexus 5010/5020 switches using Fibre cables.
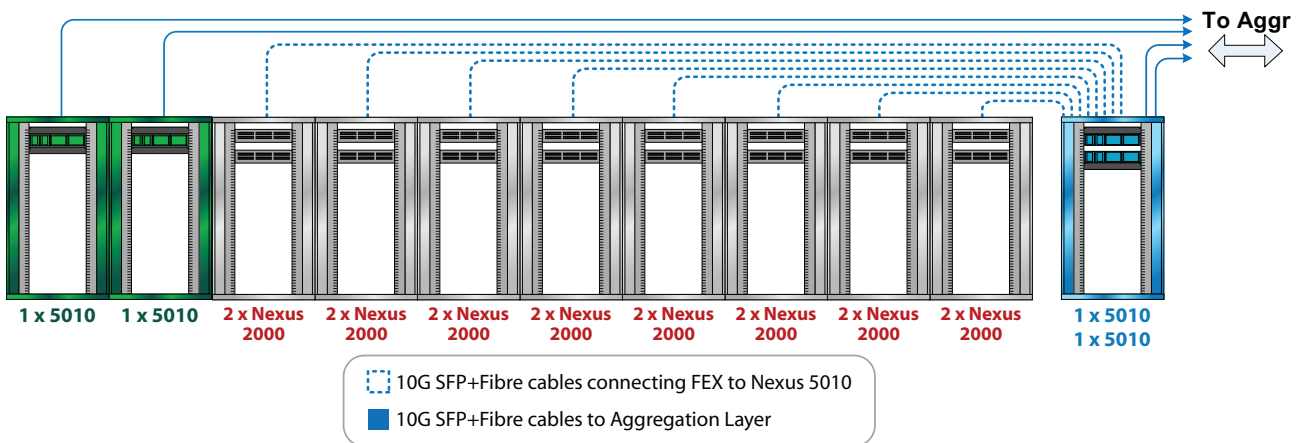

Fig. 1.8

Summary

Data centers are undergoing a fundamental shift. Application developments and server virtualization is driving the need for a more flexible, efficient and available infrastructure that can dynamically adapt to the needs of the business. To deliver upon this vision, data center architectures will change to become more modular, where DC managers can easily add racks of servers by adding racks that are pre-cabled to a top-of-rack 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 inter-switch connectivity. The transition to 10G will also enable the adoption of a unified fabric using Fibre Channel over Ethernet when organizations are ready to take advantage of the technology to drive further efficiencies within the data center.

As part of the data center evolution, cabling infrastructure must be considered in the context of delivering reliable high bandwidth and available application performance. This will also require careful consideration with respect to maximizing power and cooling efficiency, as power and cooling has emerged as the number-one issue for many organizations. Another consideration for Greenfield and Brownfield data center design is the modular approach to data cabling and network architecture (a *pod* in this paper). The modular approach provides improves data center design, deployment and operations flexibility, efficiency and responsiveness. This in turn lowers overall CAPEX and OPEX. The Top of Rack architecture fits the modularity principle and enables the "rack-and-roll" deployment model for further modularity and flexibility at the rack level. The ToR network architecture and cabling architecture model provides a graceful transition from 1G to 10G attached servers while providing infrastructure readiness to adopt 40G and 100G server and switch to switch connectivity without changing the cabling plant.