PANDUIT®

Data Centers

Mapping Cisco Nexus, Catalyst, and MDS Logical Architectures into *PANDUIT* Physical Layer Infrastructure Solutions

DESIGN GUIDE

Introduction

The growing speed and footprint of data centers is challenging IT Managers to effectively budget and develop reliable, high-performance, secure, and scalable network infrastructure environments. This growth is having a direct impact on the amount of power and cooling required to support overall data center demands. Delivering reliable power and directly cooling the sources that are consuming the majority of the power can be extremely difficult if the data center is not planned correctly.

This design guide examines how physical infrastructure designs can support a variety of network layer topologies in order to achieve a truly integrated physical layer infrastructure. By understanding the network architecture governing the arrangement of switches and servers throughout the data center, network stakeholders can map out a secure and scalable infrastructure design to support current applications and meet anticipated bandwidth requirements and transmission speeds.

The core of this guide presents a virtual walk through the data center network architecture, outlining the relationships of key physical layer elements including switches, servers, power, thermal management, racks/cabinets, cabling media, and cable management. The deployment of Cisco hardware in two different access models are addressed: Top of Rack (ToR) and End of Row (EoR).

Top of Mind Data Center Issues

The following issues are critical to the process of building and maintaining cost-effective network infrastructure solutions for data centers:

- Uptime
  Uptime is the key metric by which network reliability is measured, and can be defined as the amount of time the network is available without interruption to application environments. The most common service interruptions to the physical layer result from operational changes.

- Scalability
  When designing a data center, network designers must balance today's known scalability requirements with tomorrow's anticipated user demands. Traffic loads and bandwidth/distance requirements will continue to vary throughout the data center, which translates to a need to maximize your network investment.

- Security
  A key purpose of the data center is to house mission critical applications in a reliable, secure environment. Environmental security comes in many forms, from blocking unauthorized access to monitoring system connectivity at the physical layer. Overall, the more secure your network is, the more reliable it is.

Developing the Integrated Infrastructure Solution

Data center planning requires the close collaboration of business, IT, and facilities management teams to develop an integrated solution. Understanding some general planning relationships helps you translate business requirements into practical data center networking solutions.

Business Requirements Drive Data Center Design

A sound data center planning process is inclusive of the needs of various business units. Indeed, the process requires the close collaboration of business, IT, and facilities management teams to develop a broad yet integrated data center solution set. Understanding some general planning relationships helps you translate business requirements into practical data center networking solutions.

Business requirements ultimately drive all data center planning decisions (see Figure 1). On a practical level, these requirements directly impact the type of applications deployed and Service Level Agreements (SLAs) adopted by the organization. Once critical uptime requirements are in place and resources (servers, storage, compute, and switches) have been specified to support mission-critical applications, the required bandwidth, power, and cooling loads can be estimated.

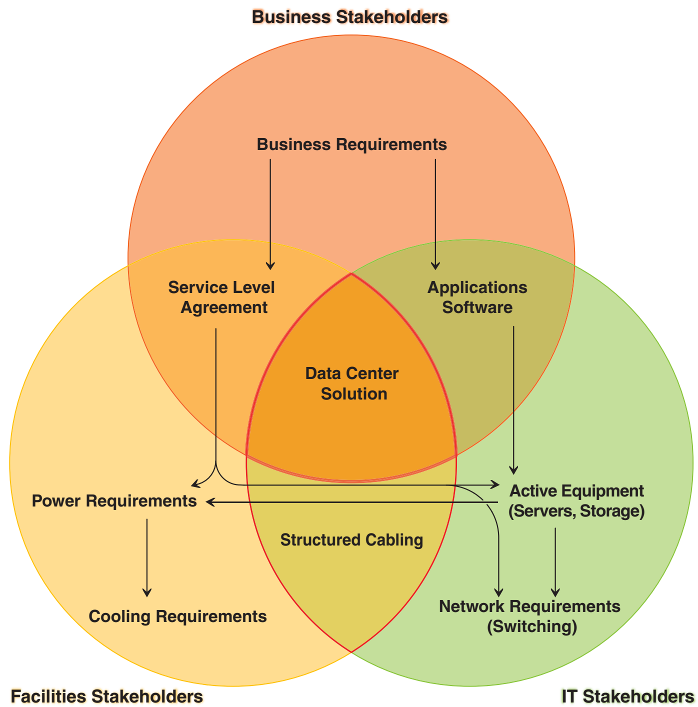

Figure 1. Business Requirements

Some data center managers try to limit the number of standard compute resources on fewer hardware platforms and operating systems ,which makes planning decisions related to cabinet, row, and room more predictable over time. Other managers base their design decisions solely on the business application, which presents more of a challenge in terms of planning for future growth. The data center network architecture discussed in this guide uses a modular, End of Row (EoR) model and includes all compute resources and their network, power, and storage needs. The resulting requirements translates to multiple LAN, SAN, and power connections at the physical layer infrastructure. These requirements can be scaled up from cabinet to row, from row to zone, and from zones to room.

Designing Scalability into the Physical Layer

When deploying large volumes of servers inside the data center it is extremely important that the design footprint is scalable. However, access models vary between each network, and can often be extremely complex to design. The integrated network topologies discussed in this guide take a modular, platform-based approach in order to scale up or down as required within a cabinet or room. It is assumed that all compute resources incorporate resilient network, power, and storage resources. This assumption translates to multiple LAN, SAN, and power connections within the physical layer infrastructure.

One way to simplify the design and simultaneously incorporate a scalable layout is to divide the raised floor space into modular, easily duplicated sub-areas. Figure 2 illustrates the modular building blocks used in order to design scalability into the network architecture at both OSI Layers 1 and 2. The logical architecture is divided into three discrete layers, and the physical infrastructure is designed and divided into manageable sub-areas called "Pods". This example shows a typical data center with two zones and 20 Pods distributed throughout the room; core and aggregation layer switches are located in each zone for redundancy, and access layer switches are located in each Pod to support the compute resources within the Pod.

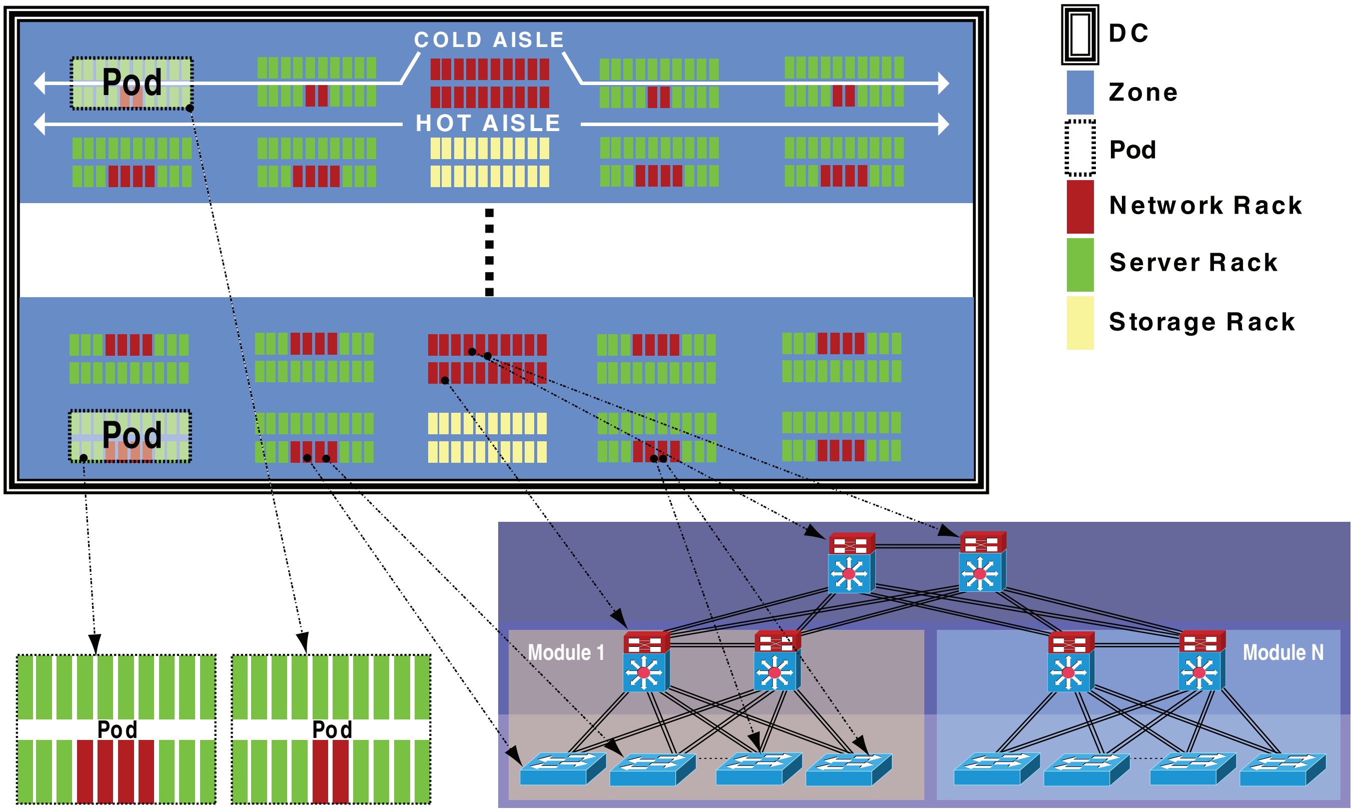

Figure 2. Mapping the Logical Architecture to the Cabling Infrastructure

Network Access Layer Environments

This guide describes two models for access layer switching environments – Top of Rack (TOR), and End of Row (EOR) – and reviews the design techniques needed for the successful deployment of these configurations within an integrated physical layer solution. When determining whether to deploy a TOR or EOR model it is important to understand the benefits and challenges associated with each:

- A ToR design reduces cabling congestion which enhances flexibility of network deployment and installation. Some trade-offs include reduced manageability and network scalability for high-density deployments due to the need to manage more access switches than in an EoR configuration.

- An EoR model (also sometimes known as Middle of Row [MoR]) leverages chassis-based technology for one or more row of servers to enable higher densities and greater scalability throughout the data center. Large modular chassis such as the Cisco Nexus 7000 Series and Cisco Catalyst 6500 Series allow for greater densities and performance with higher reliability and redundancy. Figure 2 represents an EoR deployment with multiple Pods being distributed throughout the room

Note: Integrated switching configurations, in which applications reside on blade servers that have integrated switches built into each chassis, are not covered in this guide. These designs are used only in conjunction with blade server technologies and would be deployed in a similar fashion as EoR configurations.

Top of Rack (ToR) Model

The design characteristic of a ToR model is the inclusion of an access layer switch in each server cabinet, so the physical layer solution must be designed to support the switching hardware and access-layer connections. One cabling benefit of deploying access layer switches in each server cabinet is the ability to link to the aggregation layer using long-reach small form factor fiber connectivity. The use of fiber eliminates any reach or pathway challenges presented by copper connectivity to allow greater flexibility in selecting the physical location of network equipment.

Figure 3 shows a typical logical ToR network topology, illustrating the various redundant links and distribution of connectivity between access and aggregation switches. This example utilizes the Cisco Nexus 7010 for the aggregation layer and a Cisco Catalyst 4948 for the access layer. The Cisco Catalyst 4948 provides 10GbE links routed out of the cabinet back to the aggregation layer and 1GbE links for server access connections within the cabinet.

Once the logical topology has been defined, the next step is to map a physical layer solution directly to that topology. With a ToR model it is important to understand the number of network connections needed for each server resource. The basic rule governing the number of ToR connections is that any server deployment requiring more than 48 links requires an additional access layer switch in each cabinet to support the higher link volume. For example, if thirty (30) 1 RU servers that each require three copper and two fiber connections are deployed within a 45 RU cabinet, an additional access layer switch is needed for each cabinet. Figure 4 shows the typical rear view ToR design including cabinet connectivity requirements at aggregation and access layers.

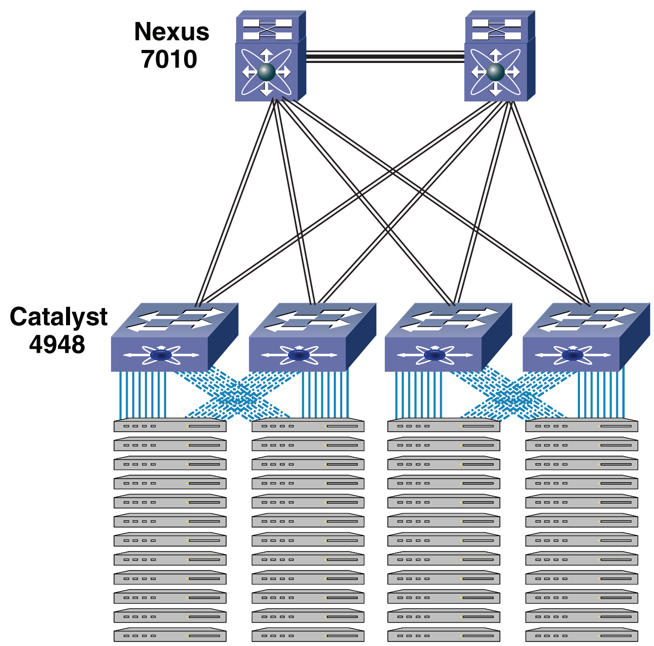

Figure 3. Logical ToR Network Topology

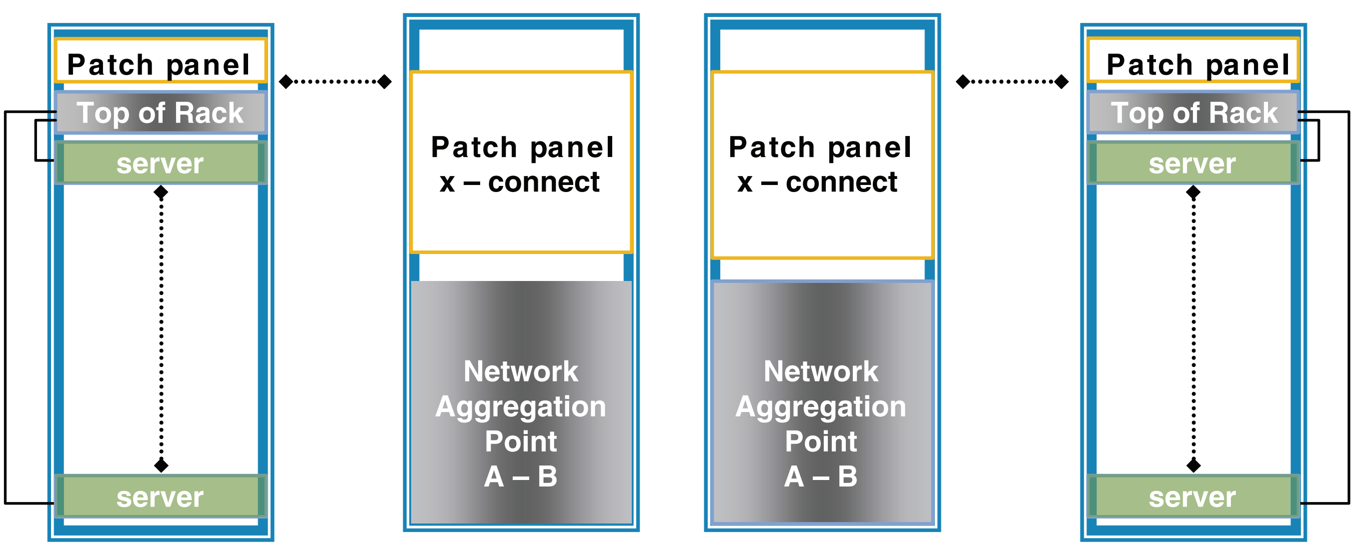

Figure 4. Rear View of ToR Configuration

Density considerations are tied to the level of redundancy deployed to support mission critical hardware throughout the data center. It is critical to choose a deployment strategy that accommodates every connection and facilitates good cable management.

High-density ToR deployments like the one shown in Figure 5 require more than 48 connections per cabinet. Two access switches are deployed in each *Net-Access*™ Cabinet to support complete redundancy throughout the network. All access connections are routed within the cabinet and all aggregation linked are routed up and out of the cabinet through the *FiberRunner*® Cable Routing System back to the horizontal distribution area.

Lower-density ToR deployments require less than 48 connections per cabinet (see Figure 6). This design shares network connections between neighbor cabinets to provide complete redundancy to each compute resource.

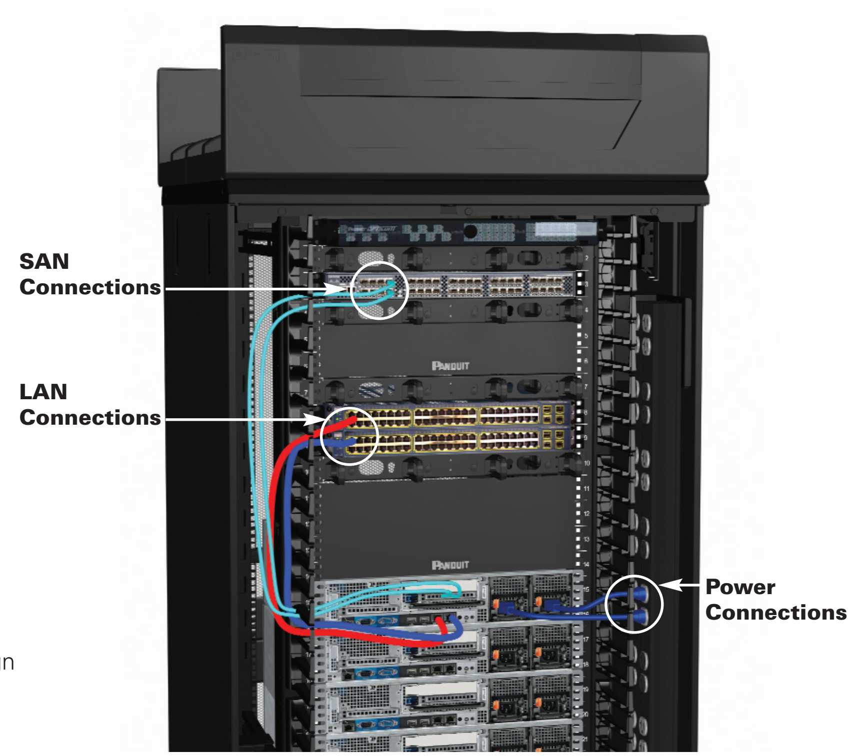

Figure 5. Dual Switch – Server Cabinet Rear View

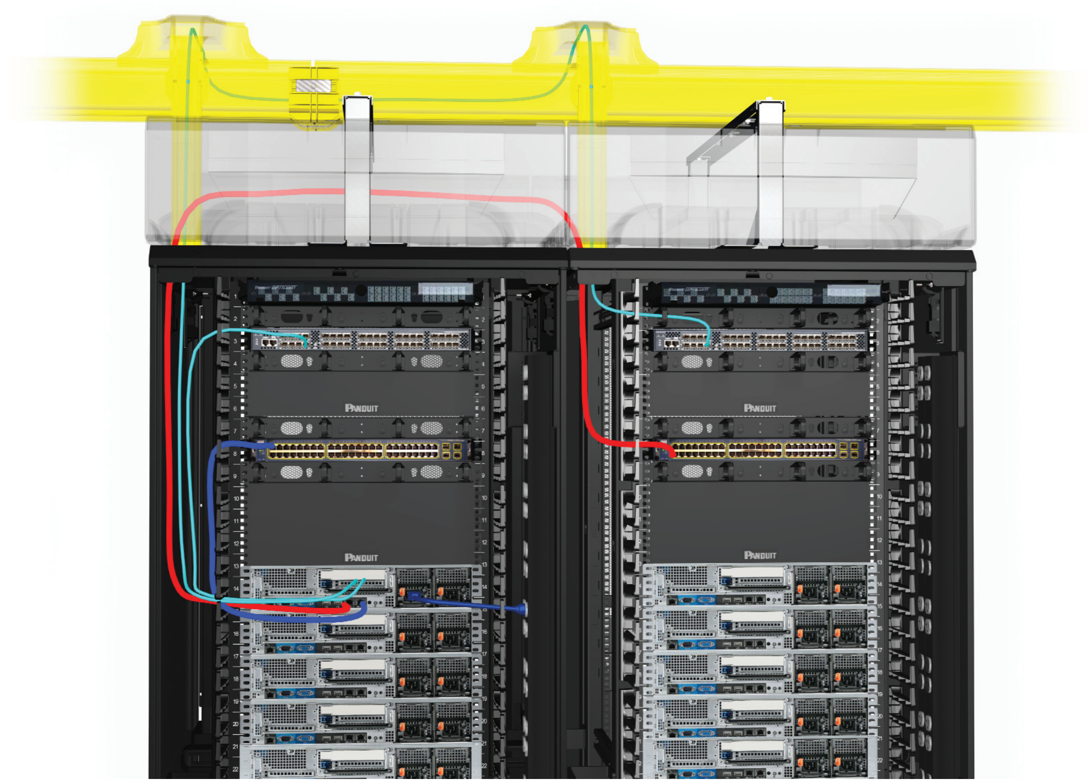

Figure 6. Single Switch – Server Cabinet Rear View

The cross-over of network connections between cabinets presents a cabling challenge that needs to be addressed in the physical design to avoid problems when performing any type of operational change after initial installation. To properly route these connections between cabinets there must be dedicated pathways defined between each cabinet to accommodate the cross-over of connections. The most common approach is to use *PANDUIT* overhead cable routing systems that attach directly to the top of the *NET-ACCESS*™ Cabinet to provide dedicated pathways for all connectivity routing between cabinets, as shown in Figure 6.

Table 1 itemizes the hardware needed to support a typical ToR deployment in a data center with 16,800 square feet of raised floor space. Typically in a ToR layout, the Cisco Catalyst 4948 is located towards the top of the cabinet. This allows for heavier equipment such as servers to occupy the lower portion of the cabinet closer to the cooling source to allow for proper thermal management of each compute resource. In this layout, 1 RU servers are specified at 24 servers per cabinet, with two LAN and two SAN connections per server to leverage 100% of the LAN switch ports allocated to each cabinet.

Connectivity is routed overhead between cabinets to minimize congestion and allow for greater redundancy within the LAN and SAN environment. All fiber links from the LAN and SAN equipment are routed via overhead pathways back to the Cisco Nexus, Catalyst, and MDS series switches at the aggregation layer. The *PANDUIT® FIBERRUNNER®* Cable Routing System supports overhead fiber cabling, and *PANDUIT® NET-ACCESS*™ Overhead Cable Routing System can be leveraged with horizontal cable managers to support copper routing and patching between cabinets.

Data Center Assumptions	Quantity	Server Specifications:
Raised Floor Square Footage	16800	2 — 2.66 GHz Intel Quad Core Xeon X5355
Servers	9792	2 x 670 W Hot-Swap
Cisco Nexus 7010	10	4 GB Ram — (2) 2048MB Dimm(s)
Cisco Catalyst 4948	408	2 — 73GB 15K-rpm Hot-Swap SAS — 3.5
Server Cabinets	408	
Network Cabinets (LAN & SAN Equipment)	33	
Midrange & SAN Equipment Cabinets	124	

	Watts Per Device	Per Cabinet	Cabs Per Pod	Pods Per Room	Quantity	Total Watts
Servers	350	24	102	4	9792	3,427,200
Cisco Catalyst 4948 Switches (Access)	350	1	102	4	408	142,800
Cisco MDS (Access)	100	1	102	4	408	40,800
Cisco Appliance Allocation	5,400	—	—	—	8	43,200
Cisco Nexus 7010 (Aggregation)	5,400	—	2	4	8	43,200
Cisco MDS (Aggregation)	5,400	—	2	4	8	43,200
Cisco Nexus 7010 (Core)	5,400	—	—	—	2	10,800
Cisco MDS (Core)	5,400	—	—	—	2	10,800
Midrange/SAN Equipment Cabinets	4,050	—	—	—	124	502,200
Midrange/SAN Switching MDS	5,400	—	—	—	4	21,600
Midrange/SAN Switching Nexus 7010	5,400	—	—	—	4	21,600
					Total Watts:	4,307,400
					Total Kilowatts:	4,307.40
					Total Megawatts:	4.31

Table 1. Room Requirements for Typical ToR Deployment

End of Row (EoR) Model

In an EoR model, server cabinets contain patch fields but not access switches. In this model, the total number of servers per cabinet and I/Os per server determines the number of switches used in each Pod, which then drives the physical layer design decisions. The typical EoR Pod contains two Cisco Nexus or Cisco Catalyst switches for redundancy. The length of each row within the Pod is determined by the density of the network switching equipment as well as the distance from the server to the switch. For example, if each server cabinet in the row utilizes 48 connections and the switch has a capacity for 336 connections, the row would have the capacity to support up to seven server cabinets with complete network redundancy, as long as the seven cabinets are within the maximum cable length to the switching equipment.

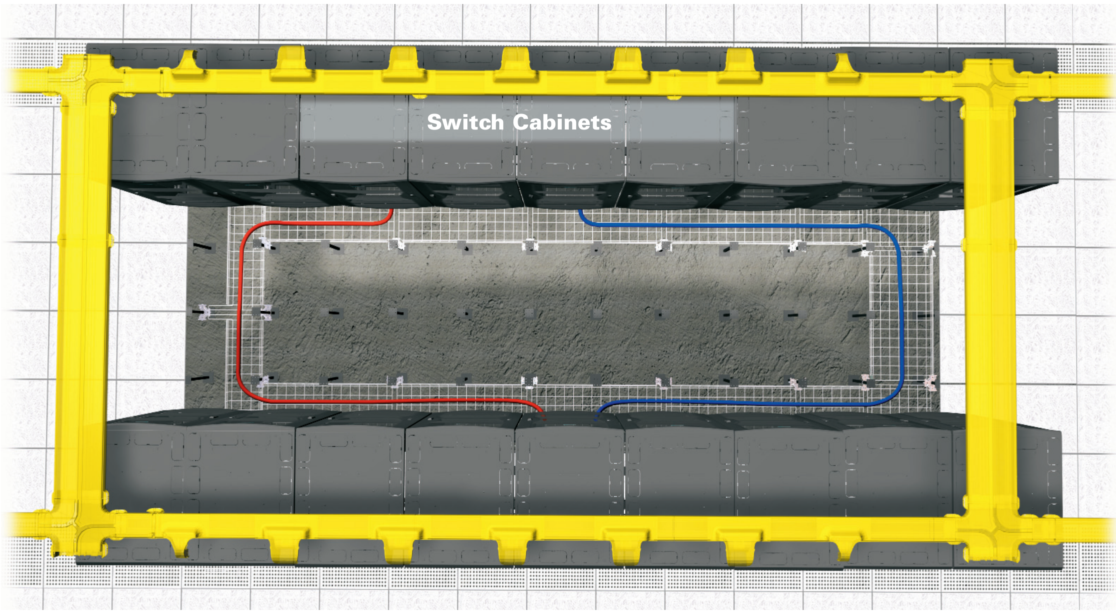

Figure 7. Top View of EoR Configuration

Figure 7 depicts a top view of a typical Pod design for a EoR configuration, and shows proper routing of connectivity from a server cabinet to both access switches. Network equipment is located in the middle of the row to distribute redundant connections across two rows of cabinets to support a total of 14 server cabinets. The red line represents copper LAN "A" connections and the blue line represents copper LAN "B" connections. For true redundancy the connectivity takes two diverse pathways using *PANDUIT® GRIDRUNNER™* Underfloor Cable Routing Systems to route cables underfloor to each cabinet. For fiber connections there is a similar pathway overhead to distribute all SAN "A" and "B" connections to each cabinet.

As applications continue to put even greater demands on the network infrastructure it is critical to have the appropriate cabling infrastructure in place to support these increased bandwidth and performance requirements. Each EoR-arranged switch cabinet is optimized to handle the high density requirements from the Cisco Nexus 7010 switch. Figure 8 depicts the front view of a Cisco Nexus 7010 switch in a *NET-ACCESS™* Cabinet populated with Category 6A 10 Gigabit cabling leveraged for deployment in an EoR configuration. It is critical to properly manage all connectivity exiting the front of each switch. The EoR-arranged server cabinet is similar to a typical ToR-arranged cabinet (see Figure 5), but the characteristic access layer switch for both LAN and SAN connections are replaced with structured cabling.

Table 2 itemizes the hardware needed to support a typical EoR deployment in a data center with 16,800 square feet of raised floor space. The room topology for the EoR deployment is not drastically different from the ToR model. The row size is determined by the typical connectivity requirements for any given row of server cabinets. Most server cabinets contain a minimum of 24 connections and sometimes exceed 48 connections per cabinet. All EoR reference architectures are based around 48 copper cables and 24 fiber strands for each server cabinet.

Figure 8. *PANDUIT® NET-ACCESS™* Cabinet with Cisco Nexus Switch

Data Center Assumptions	Quantity	Server Specifications:
Raised Floor Square Footage	16800	2 — 2.66 GHz Intel Quad Core Xeon X5355
Servers	9792	2 x 835 W Hot-Swap
Cisco Nexus 7010	50	16 GB Ram — (4)4096MB Dimm(s)
Server Cabinets	336	2 — 146GB 15K-rpm Hot-Swap SAS — 3.5
Network Cabinets (LAN & SAN Equipment)	108	
Midrange & SAN Equipment Cabinets	132	

	Watts Per Device	Per Cabinet	Cabs Per Pod	Pods Per Room	Quantity	Total Watts
Servers	450	12	14	24	9792	1,814,400
Cisco Appliance Allocation	5,400	—	—	—	8	43,200
Cisco Nexus 7010 (Access)	5,400	—	2	24	48	259,200
Cisco MDS (Access)	5,400	—	2	24	48	259,200
Cisco Nexus 7010 (Core)	5,400	—	—	—	2	10,800
Cisco MDS (Core)	5,400	—	—	—	2	10,800
Midrange/SAN Equipment Cabinets	4,050	—	—	—	132	502,200
Midrange/SAN Switching MDS	5,400	—	—	—	4	21,600
Midrange/SAN Switching Nexus 7010	5,400	—	—	—	4	21,600
					Total Watts:	4,307,400
					Total Kilowatts:	4,307.40
					Total Megawatts:	4.31

Table 2. Room Requirements for Typical EoR Deployment

Considerations Common to ToR and EoR Configurations

PANDUIT is focused on providing high-density, flexible physical layer solutions that maximize data center space utilization and optimize energy use. The following sections describe cabinet, cooling, and pathway considerations that are common to all logical architectures.

Cabinets

Cabinets must be specified that allow for maximum scalability and flexibility within the data center. The vertically mounted patch panel within the cabinet provides additional rack units that can be used to install more servers within the 45 rack units available. These vertically mounted panels also provide superior cable management versus traditional ToR horizontal patch panels by moving each network connection closer to the server network interface card ultimately allowing a shorter patching distance with consistent lengths throughout the cabinet.

Figure 9 represents a typical vertical patch panel deployment in the *NET-ACCESS*™ Server Cabinet. Complete power and data separation is achieved through the use of vertical cable management on both sides of the cabinet frame. The cabinet also allows for both overhead and underfloor cable routing for different data center applications. Shorter power cords are an option to remove the amount of added cable slack from longer cords that are shipped with server hardware. Using shorter patch cords and leveraging *PANDUIT* vertical cable management integrated into the *NET-ACCESS*™ Cabinet alleviates potential airflow issues at the back of each server that could result from poorly managed cabling.

Figure 9. Cabinet Vertical Mount Patch Panel

Thermal Management

Data center power requirements continue to increase at high rates making it difficult to plan appropriately for the proper cooling systems needed to support your room. Cabinets play a critical role in managing the high heat loads generated by active equipment. Each cabinet will require different power loads based upon the type of servers being installed as well as the workload being requested of each compute resource. Understanding cabinet level power requirements gives greater visibility into overall room conditions.

PANDUIT Laboratories' research into thermal management includes advanced computational fluid dynamics (CFD) analysis to model optimal airflow patterns and above-floor temperature distributions throughout the data center. This data is then used to develop rack, cabinet, and cable management systems that efficiently route and organize critical IT infrastructure elements. Figure 10 represents an analysis done based upon the assumptions for the 16,800 square foot EoR deployment.

PANDUIT® NET-ACCESS™ Cabinets feature large pathways for efficient cable routing and improved airflow while providing open-rack accessibility to manage, protect, and showcase cabling and equipment. Elements such as exhaust ducting, filler panels, and the *PANDUIT® COOL BOOT™* Raised Floor Assembly support hot and cold aisle separation in accordance with the TIA-942 standard. These passive solutions (no additional fans or compressors) contribute by minimizing bypass air in order to manage higher heat loads in the data center and ensure proper equipment operation and uptime.

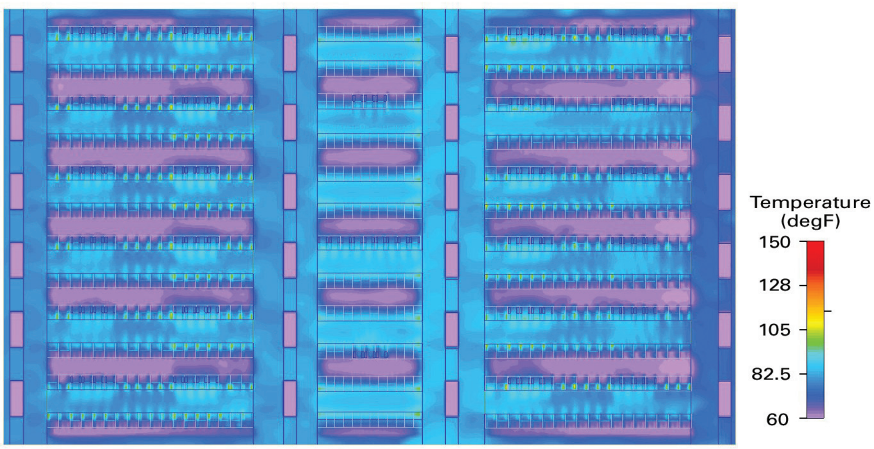

Designed in a Hot – Cold Architecture

12 – 20 Ton CRAC Units Outside

12 – 30 Ton CRAC Units Inside

Utilizing Ceiling Plenum for Return Air

All Perforated Tiles at 25% Open

Peak Temp was 114°

Figure 10. Data Center EoR Thermal Analysis

Pathways

The variety and density of data center cables means that there is no "one size fits all" solution when planning cable pathways. Designers usually specify a combination of pathway options. Many types and sizes are available for designers to choose from, including wire basket, ladder rack, J-hooks, conduit, solid metal tray, and fiber-optic cable routing systems. Factors such as room height, equipment cable entry holes, rack and cabinet density, and cable types, counts, and diameters also influence pathway decisions. The pathway strategies developed for ToR and EoR models all leverage the *FIBERRUNNER®* Cable Routing System to route horizontal fiber cables, and use the *NET-ACCESS™* Overhead Cable Routing System in conjunction with a wire basket or ladder rack for horizontal copper and backbone fiber cables.

This strategy offers several benefits:

- The combination of overhead fiber routing system and cabinet routing system ensures physical separation between the copper and fiber cables, as recommended in the TIA-942 data center standard

- Overhead pathways such as the *PANDUIT® FIBERRUNNER®* Cable Routing System protect fiber optic jumpers, ribbon interconnect cords, and multi-fiber cables in a solid, enclosed channel that provides bend radius control, and the location of the pathway is not disruptive to raised floor cooling

- The overall visual effect is organized, sturdy, and impressive

Conclusion

Next-generation data center hardware such as the Cisco Nexus 7010 switch provides increased network capacity and functionality in the data center, which in turn is placing greater demands on the cabling infrastructure. This guide describes the ways that *PANDUIT* structured cabling solutions map easily to the logical architectures being deployed in today's high-performance networks to achieve a unified physical layer infrastructure. Network stakeholders can use modular designs for both hardware architectures and cabling layouts to ensure that the system will scale over the life of the data center to survive multiple equipment refreshes and meet aggressive uptime goals.

About *PANDUIT*

PANDUIT is a leading, world-class developer and provider of innovative networking and electrical solutions. For more than 50 years, *PANDUIT* has engineered and manufactured end-to-end solutions that assist our customers in the deployment of the latest technologies. Our global expertise and strong industry relationships make *PANDUIT* a valuable and trusted partner dedicated to delivering technology-driven solutions and unmatched service. Through our commitment to innovation, quality and service, *PANDUIT* creates competitive advantages to earn customer preference.

Cisco, Cisco Systems, Nexus, Catalyst, and MDS are registered trademarks of Cisco Technology, Inc.

www.panduit.com • cs@panduit.com • 800-777-3300

PANDUIT®

Printed in U.S.A.
WW-CPCB36
8/2008