

Data Center Architecture with Panduit, Intel, and Cisco



10GBASE-T Application Note

Integrating Panduit Category 6A Interconnects with the Cisco Nexus 2232™ and Intel Ethernet Server Adapter X520-T2

Introduction

Modern data centers face a wide array of challenges that threaten the integrity of mission critical business applications. These data centers are now stretched to capacity in terms of power, cooling, and floor space. Increasing use of server consolidation, unified networking and virtualization technologies offer cost savings, improved performance, and reliability, but are driving unprecedented demand for network I/O. While Gigabit Ethernet (GbE) connections might be adequate to handle the bandwidth requirements of a single physical server, they do not have adequate bandwidth for multiple virtual servers and multiple traffic types during peak periods. The practice of adding additional GbE ports to servers can help with bandwidth, but the additional complexity, cost and administrative overhead is not sustainable. In addition, load-balancing traffic among multiple GbE links might lead to asymmetric usage of links contributing to system inefficiencies.

This is why today's IT executives and network administrators look to 10 Gigabit Ethernet (10 GbE) as the right solution to meet their escalating I/O requirements and help simplify their networks. While there are several methods for achieving a 10 GbE network, what is needed is a cost effective technology that allows broad deployment of 10 GbE to every rack and tower in the data center and backward compatibility with existing equipment. This application note examines 10GBASE-T, one of the most promising solutions for 10 GbE networks, and some of the new products for implementing it.

10GBASE-T

10GBASE-T is the most recent version of the BASE-T (commonly called Twisted Pair Ethernet) networking standard. It uses Category 6A or higher balanced twisted pair copper cabling for a maximum reach of 100 meters (328 feet). As defined by IEEE 802.3an, 10GBASE-T uses the same RJ45 connector interface found in most networks today. Preferred because of its low cost and flexibility, BASE-T networking has been the cabling infrastructure of choice in data centers for the past 25 years.

An end-to-end 10GBASE-T infrastructure is able to support the considerable I/O demands of virtualized servers and I/O intensive applications while providing high availability, agility, and security for the data center. Further, 10GBASE-T allows for auto-negotiation, enabling equipment to work with previous generations of BASE-T Ethernet. This backward compatibility facilitates seamless migration from GbE to 10 GbE, allowing for phased upgrades as desired.

10GBASE-T Interconnect Advantages

Data Center administrators have a number of 10 GbE interconnect options: 10GBASE-SR (Short Reach fiber optic), Enhanced Small Form-Factor Pluggable SFP+ Direct Attach Copper (DAC), and 10GBASE-T. Currently the most commonly used options are 10GBASE-SR optical and SFP+ DAC copper because they offer low latency and can both be used with the same SFP+ ports. However, neither 10GBASE-SR optical nor SFP+ DAC are compatible with existing installed structured cabling or with lower speed equipment such as 1000BASE-T switches, making upgrades more complex and expensive.

Now, with the introduction of 10GBASE-T enabled equipment, such as the Cisco Nexus® 2232™ 10GE Fabric Extender (FEX) and the Intel Ethernet Server Adapter X520-T2, data center managers can start enjoying the many advantages of 10GBASE-T.

- The simplicity of 1000BASE-T, but 10X faster
- Field terminable twisted pair cable with a familiar RJ-45 interface
- Ability to use existing knowledge and training on previous BASE-T
- Lengths up to 100 meters – sufficient reach for most data centers
- Flexible 10 GbE architectures to support all data center topologies
- Auto-negotiable speed offering backward compatibility with slower BASE-T networks allowing phased upgrades for a seamless migration to 10 GbE
- Lowest cost of deployment
- Ability to integrate server port LAN on motherboard (LOM) which frees up an I/O slot

10GBASE-T Power Consumption Is Reduced

One commonly held misconception of 10GBASE-T is a concern over high power consumption. It should be noted that this was also true of early 1000BASE-T PHYs, but continual advances in manufacturing processes now have the 1000BASE-T PHYs below 1W per port. 10GBASE-T adapter cards are expected to follow this same trend. 10GBASE-T power consumption has been rapidly dropping and now Intel's third-generation 10GBASE-T adapter card, the dual-port Intel Ethernet Server Adapter X520-T2, which includes both Media Access Controller (MAC) and PHY, uses less than 10W per port. This is less than half the power required by Intel's first-generation single-port product. Intel's upcoming fourth-generation dual-port 10GBASE-T controller will integrate the MAC and PHY into a single package, which will again drop power consumption significantly per port. Future applications will also look to employ Energy Efficient Ethernet (EEE), now available on GbE adapters, to further lower power consumption. These reductions in power, combined with steady performance increases, make 10GBASE-T suitable for high-density switches and enable LOM implementations on mainstream servers, further reducing the total system cost.


10GBASE-T Implementation Overview

In order to deliver a full end-to-end 10GBASE-T solution, all network elements, such as cabling, server adapters, and switching fabric, must support this higher speed. Panduit, Intel, and Cisco each deliver technologies that together form a comprehensive, integrated 10GBASE-T ecosystem solution that has been tested to ensure interoperability, availability, and reliability for 10 GbE high speed networks. This cost-effective and flexible system can be developed through the integration of the new 10GBASE-T enabled Cisco Nexus 2232™ 10 GbE fabric extenders, working with Cisco Nexus 5000 Series switches and Intel Ethernet Server Adapters over the Panduit TX6A™ 10Gig™ Copper Cabling System.

This solution can be deployed in Top-of-Rack (ToR), Middle-of-Row (MoR), End-of-Row (EoR), and distributed topologies. All of these physical topologies provide modularity, organization, and flexibility within the data center. Each of these topologies offers different advantages for both new Data Centers and existing Data Centers migrating from GbE to 10 GbE.

Panduit® TX6A™ and TX6A-SD™ 10Gig™ Copper Cabling

Other advances in 10GBASE-T have come with new higher performance and smaller diameter copper cabling. Early generations of Category 6A 10 GbE compliant cabling were rather large in diameter, up to 0.350 inch, in order to suppress cable to cable coupling known as alien crosstalk. Newer cabling technology such as Panduit's TX6A™ 10Gig™ UTP (Unshielded Twisted Pair) cabling with patented MaTriX Technology has reduced the cable diameter to a more manageable 0.300 inch while not sacrificing performance. In addition, for applications such as large data centers in which reduced cabling density is important, Panduit offers TX6A-SD™ 10Gig™ (small diameter) UTP cable with MaTriX Technology. The TX6A-SD™ 10Gig™ cable provides fully compliant Category 6A/10Gig™ performance over a channel up to 70 meters (230 ft.) with a cable diameter of only 0.240 inch. This is similar in size to previous generations of Category 6 cabling and allows full re-use of existing pathways and cable managers to facilitate seamless migration from GbE to 10 GbE.



Industry Category 6A

Panduit TX6A™

Panduit TX6A-SD™


Comparison of 100 cables in the same pathway

Cisco Nexus 2232TM 10GbE Fabric Extender

The Cisco Nexus 2232TM fabric extenders have 32 ports of 1/10GBASE-T and 8 SFP+ ports and behave as remote I/O modules for a parent Cisco Nexus 5000 Series or 7000 Series Switch. The fabric extender, essentially an extension of the Nexus switch fabric, works together with the parent switch to form a distributed modular system. The Cisco Nexus 2232TM forwards all traffic to the parent switch over 10 GbE up-links. Passing all traffic to the parent switch allows traffic to be shaped according to homogeneous policies established on the parent switch with a single point of management across more than (1,500) 1GbE ports and more than (1000) 10 GbE ports.

The distributed modular system architecture includes the following benefits:


- Architectural flexibility
- Highly scalable server access
- Simplified operations
- Increased business benefits



Cisco Nexus 2232TM

Intel Ethernet Server Adapter X520-T2

The Intel Ethernet Server Adapter X520-T2 delivers easy-to-deploy, high-volume 10 GbE connectivity over existing copper infrastructure. This card includes Intel Virtual Machine Device Queues for optimized I/O virtualization, supports unified data and storage networking, and provides reliable, consistent performance across a broad range of virtualized and non-virtualized operating systems. The Intel Ethernet Server Adapter X520-T2 delivers advantages in cost, power, and server slot savings compared to the use of multiple GbE adapters. Support for unified networking further increases consolidation savings by eliminating the need for dedicated storage adapters. Backward compatibility with 1000BASE-T infrastructures helps ensure easy deployment and upgrading.



Intel Ethernet Server Adapter X520-T2

The Intel Ethernet Server Adapter X520-T2 is Intel's third-generation 10GBASE-T adapter and includes new features to support high-volume 10GbE.

- Dual-port design for redundancy and greater throughput
- Ability to send and receive off-loads to accelerate iSCSI traffic, and optimizations for multi-core processors to improve efficiency
- Intel Virtualization Technology for Connectivity; virtual machine device queues and single-root I/O virtualization (SR-IOV) support boost throughput and reduce latency to improve overall system performance
- PCI Express Generation 2-compliance for dual-port, line-rate 10 GbE throughput

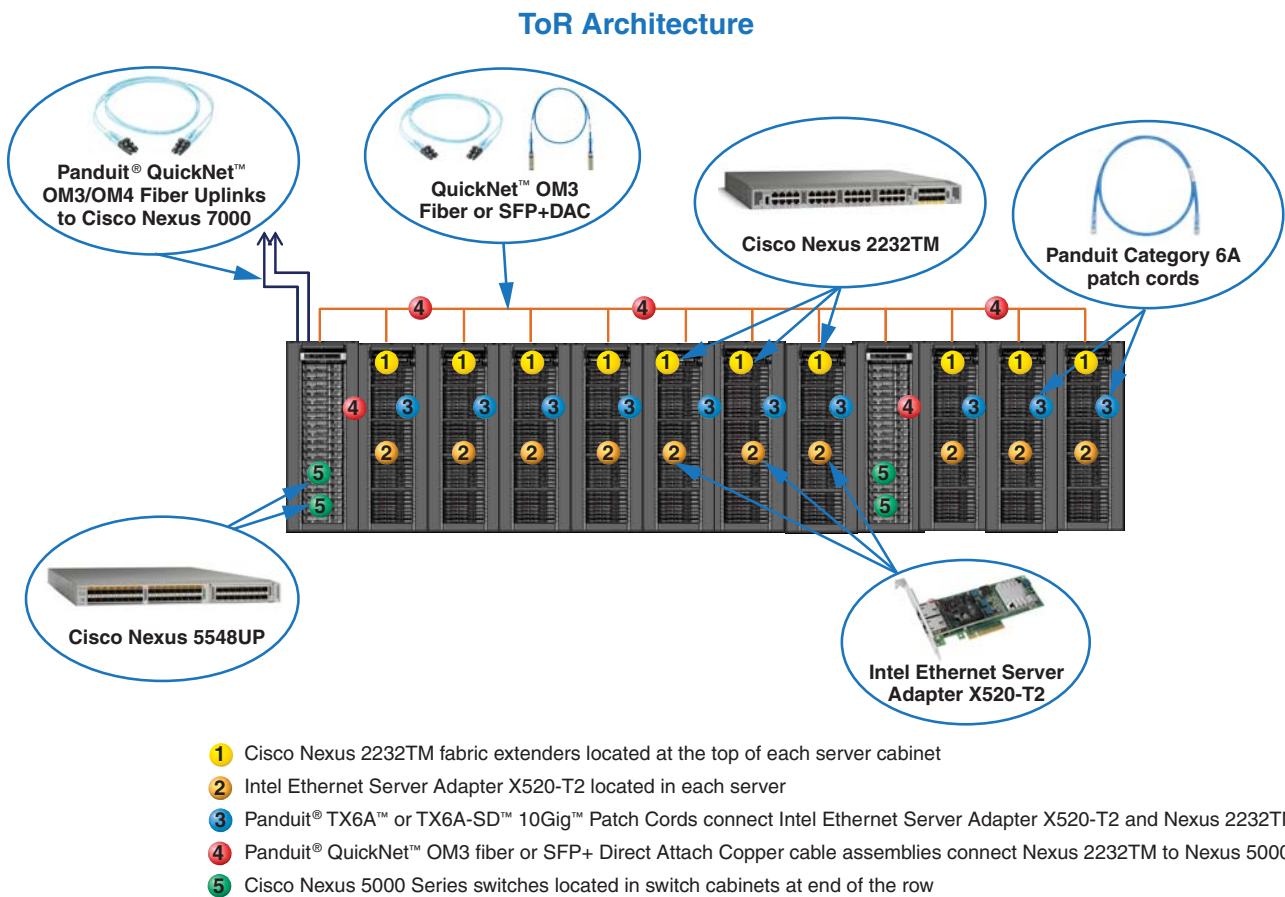
10GBASE-T Network Architecture

The new 10GBASE-T enabled Cisco Nexus 2232™ fabric extenders bring architectural flexibility allowing for support of all topologies commonly used in Data Centers including ToR, EoR, and the more traditional MoR. Each architecture has its own benefits and constraints, so the choice is dependent on the preferences of the data center manager, the existing facilities, equipment layout, and the applications being run. In the following sections we will review the advantages of each architecture and how 10GBASE-T can be efficiently implemented.

ToR Architecture

Top-of-Rack is a distributed modular system in which a switch or a fabric extender such as the Cisco Nexus 2232TM resides at the top of the server rack. Although the EoR and MoR designs are most common today, ToR is becoming popular because it allows the use of higher density server racks. As the number of cables exiting the rack increases with the number of servers the cost and space constraints become harder to maintain. When used in ToR architecture the Nexus 2232TM fabric extender aggregates network connections in the rack, reducing the number of intra-enclosure cable runs back to the Cisco Nexus 5000 Series switches.

Each Nexus 2232TM has (32) 10 GbE downlink ports, allowing it to support up to 32 servers. A common configuration has twelve cabinets per row, each containing up to sixteen 2U servers, or 32 1U servers each with a dual-port Intel Ethernet Server Adapter X520-T2. One or two Nexus 2232TM fabric extenders are placed at the top of each cabinet. The second fabric extender can be used to provide either redundancy or the capability to add more servers per cabinet.



In the above scenario, the Nexus 2232TM downlinks to the Intel Ethernet Server Adapter X520-T2 in the server cabinet via Panduit® TX6A™ 10Gig™ or TX6A-SD™ 10Gig™ Patch Cords. The Nexus 2232TM fabric extender acts as a remote line card for a Cisco Nexus 5000 Series switch located in a switch cabinet either at the end of the row or in the middle of the row. For redundancy the Nexus 5000 Series Switch are generally located in two different cabinets. The Nexus 2232TM is up-linked to the Nexus 5000 Series Switch using Panduit OM3 structured fiber or possibly SFP+ DAC, depending on the location and distance of the switch. Locating the Nexus 5000 in the middle of the row reduces cable lengths to the Nexus 2232TM and allows the use of SFP+ DAC cabling, which has a maximum reach of 7 meters (23 feet). When fiber is used, the OM3 cable is factory terminated with duplex LC connectors and is routed in overhead pathways such as the Panduit® FiberRunner® Cable Routing System.

There are eight SFP+ uplink ports on the Nexus 2232TM, for a 4:1 oversubscription ratio. To provide more redundancy, four up-links can be connected to one Nexus 5000 Series switch with the remaining four uplinks connecting to a different Nexus 5000 Series switch located in either in the same cabinet or different cabinets. One Nexus 5000 Series switch can support up to 24 Nexus 2232TM fabric extenders, but for redundancy a factor of twelve is normally used. Therefore, a typical row with 10 to 12 server cabinets, each with (32) 1U servers and two Nexus 2232TM fabric extenders would have a switch cabinet with two Nexus 5000 Series switches. Similarly, a row of eighteen server cabinets each with (32) 1U servers and two Nexus 2232TM fabric extenders would have a switch cabinet with four Nexus 5000 Series switches or two switch cabinets each with two Nexus 5000 Series switch.

Data centers with the following requirements are well suited to the ToR topology:

- New data centers without any existing cabling or equipment
- Need for scalability and ease of adding cabinets
- High-density layouts using high numbers of 1U servers per cabinet
- Reduced cable congestion
- Use with modular pre-configured cabinets

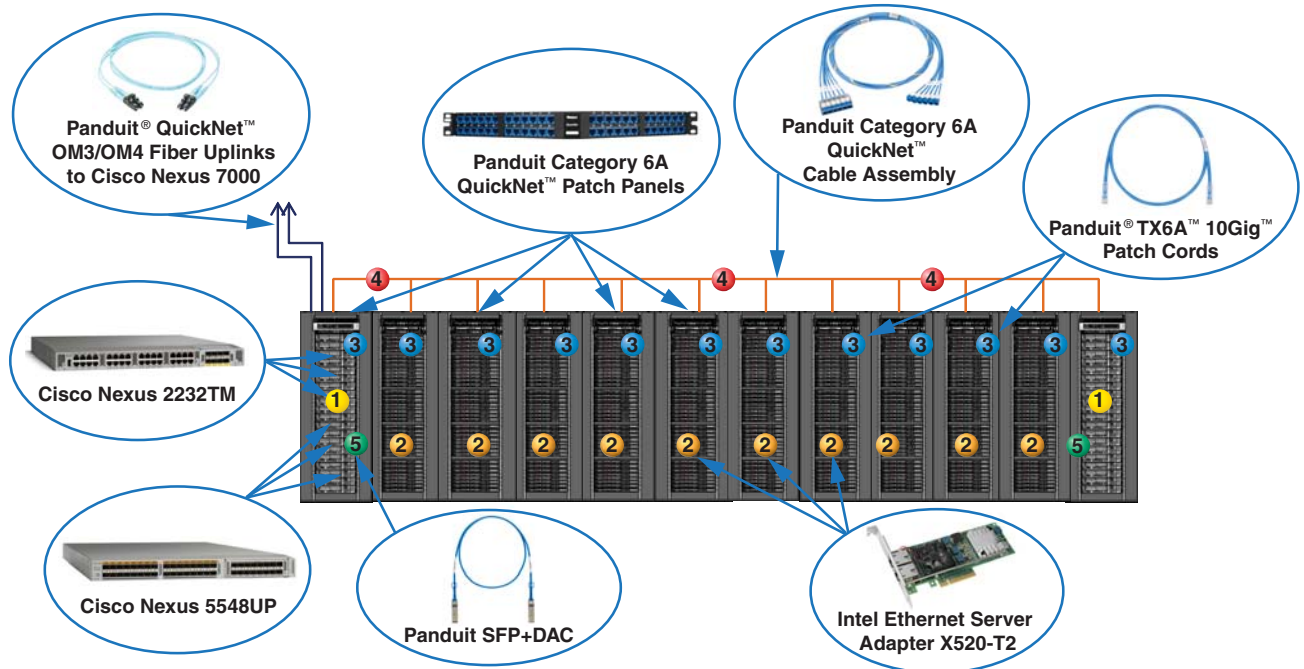
EoR Architecture

In an End-of-Row model, standard structured cabling practices are used to provide efficient cable organization. This increases the ease of maintenance and helps reduce the chances of air flow blockage. The Nexus 2232TM and Nexus 5000 Series switches are located together in switch cabinets at either end of the row and are normally mirror images of each other for redundancy. A typical EoR switch cabinet contains a quantity of Nexus 2232TM proportional to the number of servers in the row. As with ToR, a ratio of one Nexus 2232TM fabric extender for every 32 servers, or one for every sixteen servers for redundancy, is used. Like ToR, one Nexus 5000 Series switch is used per 24 (or twelve for redundancy) Nexus 2232TM fabric extenders. Using the lower ratios for redundancy, a row of eighteen server cabinets with (32) 1U servers each would typically have two switch cabinets on either end of the row, both having eighteen Nexus 2232TM fabric extenders and two Nexus 5000 Series switches. A row with twelve server cabinets would have either two switch cabinets, both with twelve Nexus 2232TM and fabric extenders one Nexus 5000 Series or one switch cabinet vertically partitioned with twelve Nexus 2232TM fabric extenders and one Nexus 5000 Series per partition. The switch cabinets or a cabinet adjacent to the switch cabinet will generally contain multiple patch panels, called patch fields, for connection to all of the server cabinets.

The server cabinets in the EoR architecture contain patch panels at the tops of the cabinets instead of switches or fabric extenders. Category 6A patch cords are used to connect the patch panels to the dual-port Intel Ethernet Server Adapter X520-T2 installed in the servers. Cabinet-to-cabinet cabling within the row connects the patch panels in the server cabinets with the switch cabinet patch field and is generally twisted pair copper cabling such as Panduit® TX6A™ 10Gig™ or TX6A-SD™ 10Gig™. To minimize congestion, the Panduit® Net-Access™ CabRunner® Overhead Cable Routing System can be used with horizontal cable managers.

End of Row architectures provide ease of replication and allow simple migration from legacy modular server access architectures. EoR allows for central patching, which reduces oversubscription and increases full port utilization. In addition, because all the switches are in just one or two cabinets, EoR simplifies network management and consolidates security into fewer locations.

EoR Architecture




- 1 Nexus 2232TM and Nexus 5000 series switches located in switch cabinets at both ends of the row
- 2 Intel Ethernet Server Adapter X520-T2 located in each server
- 3 Panduit® TX6A™ or TX6A-SD™ 10Gig™ Patch Cords connect Intel Ethernet Server Adapter X520-T2 and Nexus 2232TM to Panduit patch panels at the top of each rack
- 4 Panduit® TX6A™ or TX6A-SD™ 10Gig™ Cable and Jack Modules or QuickNet™ Pre-terminated Cable Assemblies connect the patch panels in server cabinet to patch panels in switch cabinet
- 5 Panduit SFP+ passive Direct Attach Copper cables connect Nexus 2232TM to Nexus 5000

Data centers with the following requirements are well suited to the EoR topology:

- Existing EoR/MoR networks with BASE-T cabling – currently the most common layout
- Compatibility with existing GbE or slower networks
- Potential for frequent network changes, which are easier to do with structured cabling
- Centralized security and network management
- Sensitivity to down time and strong value on redundancy
- Optimized port utilization

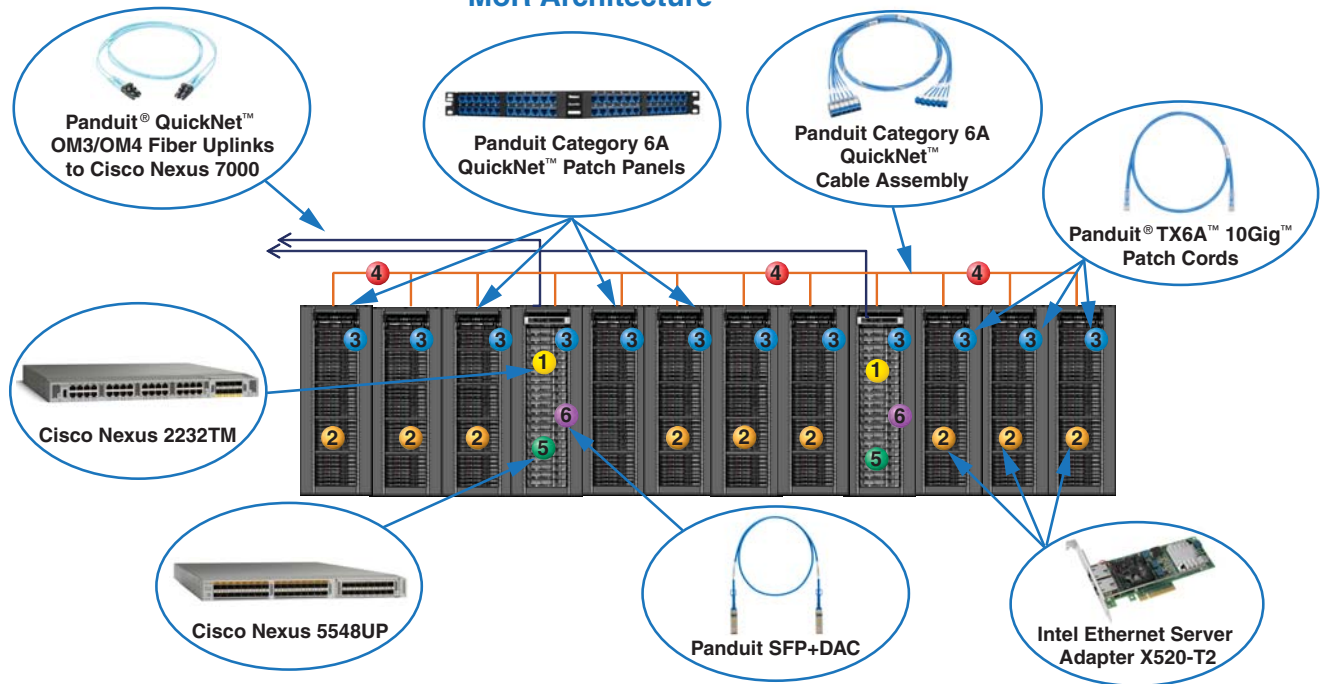
MoR Architecture

Middle of Row architecture is basically the same as EoR architecture except that the switch cabinets are located in the middle of the row, with cabling branching to the right and left to each of the server cabinets. The cable lengths are generally shorter, which can reduce cabling complexity and allows more connectivity options. The shorter runs can use lower cost copper instead of fiber and allow for single patch fields instead of the dual point patching normally used in EoR. This feature makes MoR one of the easiest topologies to maintain. However, centralizing the switching cabinets can also cause power and cooling issues if not properly planned for during the layout phase. The MoR topology is well suited to pod architectures like the below illustration which provide good redundancy and compartmentalized or/segregated computing commonly required by co-location sites and larger data centers.



Top view of a typical EoR / MoR layout showing redundant cabling.

MoR Architecture



- 1 Nexus 2232TM Fabric Extenders located in switch cabinets in the middle of the row
- 2 Intel Ethernet Server Adapter X520-T2 located in each server
- 3 Panduit TX6A™ or TX6A-SD™ 10Gig™ Patch Cords connect Intel Intel Ethernet Server Adapter X520-T2 and Nexus 2232TM to Panduit patch panels at the top of each rack
- 4 Panduit TX6A™ or TX6A-SD™ 10Gig™ Cable and Jack Modules or QuickNet™ Pre-terminated Cable Assemblies connect the patch panels in server cabinet to patch panels in switch cabinet
- 5 Nexus 5000 switches located in switch cabinets in the middle of the row
- 6 Panduit SFP+ Passive Direct Attach Copper Cable Assemblies connect Nexus 2232TM to Nexus 5000

Benefits of 10GBASE-T Structured Cabling

Most data centers currently use structured cabling because it provides an efficient and cost effective solution for managing cabling while providing multiple rerouting points to handle changes or replacement in cases of equipment failure. In the EoR and MoR architectures, Panduit TX6A™ 10Gig™ Patch Cords connect the Intel Ethernet Server Adapters to Category 6A patch panels at the top of the cabinet. The patch cords are cable assemblies made with Category 6A cabling and a male plug on each end. The patch panels, such as the Panduit QuickNet™ or DP6A™ 10 Gig™ Patch Panels, come with 24, 32, or 48 10 GbE ports to match the number of servers and are used as a connection point in the cabinet to make routing changes quick and simple. Category 6A cable terminated with Category 6A jack modules on both ends connects the patch panel in the server cabinet to the patch field in the switch cabinet. This cable assembly is called horizontal cabling or “permanent length” and has a maximum reach of 90 meters in a 10GBASE-T network. The maximum 10GBASE-T channel length of 100 meters is obtained by adding up to 10 meters of Category 6A patch cords. The horizontal cable can be terminated during installation if precise lengths are needed or come as a pre-terminated cable assembly such as Panduit’s QuickNet™ Cabling System. The patch panel in the switch cabinet is then connected to the Cisco Nexus 2232TM fabric extender via another Category 6A patch cord.

Conclusion

10GBASE-T with its ability to be used with all standard data center topologies (ToR, EoR, and MoR) is well suited for widespread usage in today's data centers. It simplifies data center infrastructures by providing the highest architectural flexibility for new data centers and the easiest upgrade path for existing data centers. By offering both sufficient reach and bandwidth, 10GBASE-T meets the demands of virtualized and ever flattening networks. Building on its tried and true BASE-T legacy, 10GBASE-T will also offer the lowest cost of deployment as adoption rates increase, power consumption continues to fall and LOM becomes widespread. In summary, the combination of reliability, flexibility, backward compatibility, cost-effectiveness, ease of use and a 100 meter reach makes 10GBASE-T an ideal choice for next generation networks.

About Cisco

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

About Intel

Intel, the world leader in silicon innovation, develops technologies, products, and initiatives to continually advance how people work and live. Founded in 1968 to build semiconductor memory products, Intel introduced the world's first microprocessor in 1971. Intel is the industry leader in Ethernet adapters and has delivered quality Ethernet products for 30 years.

About Panduit

Panduit is a world-class developer and provider of leading-edge solutions that help customers optimize the physical infrastructure through simplification, increased agility and operational efficiency. Panduit's Unified Physical InfrastructureSM (UPI) based solutions give Enterprises the capabilities to connect, manage, and automate communications, computing, power, control, and security systems for a smarter, unified business foundation. Panduit provides flexible, end-to-end solutions tailored by application and industry to drive performance, operational, and financial advantages. Panduit's global manufacturing, logistics, and e-commerce capabilities, along with a global network of distribution partners, help customers reduce supply chain risk. Strong technology relationships with industry leading systems vendors and an engaged partner ecosystem of consultants, integrators, and contractors together with its global staff and unmatched service and support make Panduit a valuable and trusted partner.

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

WW-ZCAN02 10/2011

Copyright and Trademark Information

Cisco and Cisco Systems are registered trademarks of Cisco Technology, Inc.

NOTE: The information contained herein is intended as a guide for use by persons having technical skill at their own discretion and risk. Panduit disclaims any liability arising from any information contained herein or for the absence of same.