

Deploying 10GBASE-T as the Low-Cost Path to the Cloud



Integrating Panduit Category 6A Interconnects with Cisco Nexus® Switches and Intel® Ethernet Converged Network Adapters

Introduction

The increasing deployments of server virtualization, unified data and storage networking, and cloud computing and big data solutions offer cost savings, improved performance, and reliability, but these technologies are also driving unprecedented demand for network throughput. While Gigabit Ethernet (GbE) connections can handle the bandwidth requirements of a single physical server, they are inadequate to support virtualized server-consolidation scenarios or multiple traffic types during peak periods. Scaling the network by adding additional GbE ports to servers can increase bandwidth, but the additional complexity, cost, and administrative overhead are untenable. In addition, load-balancing traffic among multiple GbE links can lead to asymmetric usage of those links, contributing to system and infrastructure inefficiencies.

To address those shortcomings, many IT executives and network administrators are now deploying 10 Gigabit Ethernet (10 GbE) to meet escalating I/O requirements and help simplify both cabling infrastructure and network operations. With 10GBASE-T connections now available as integrated or modular LAN on motherboard (LOM) connections or add-in adapters, the technology provides IT organizations with a cost-effective 10 GbE media option that allows for broad deployment of 10 GbE to every rack and tower in the data center. What's more, all of this is possible while maintaining backward-compatibility with existing network equipment.

This paper examines 10GBASE-T, one of the most promising solutions for 10 GbE networks, with a focus on new products and common models for deployment in the data center.

10GBASE-T Overview

Defined by IEEE 802.3, 10GBASE-T is the most recent version of the BASE-T (commonly referred to as "twisted-pair" Ethernet) networking standard. 10GBASE-T uses Category 6A or higher balanced twisted-pair copper cabling for a maximum reach of 100 meters (~328 feet). Important for backward compatibility, 10GBASE-T uses the same RJ-45 physical interface found in all existing BASE-T networks. Preferred for its low cost and flexibility, twisted-pair with RJ-45 connectors has been the most widely used cabling in data centers for the past 25 years. 10GBASE-T also preserves the user experience and the training investment IT organizations have already made through multiple generations of BASE-T use.

An end-to-end 10GBASE-T infrastructure can support the considerable demands of virtualized servers and I/O-intensive applications, while providing high availability, agility, and security for the data center. Further, 10GBASE-T can auto-negotiate to lower speeds, enabling backwards compatibility with existing 10/100/1000 megabits-per-second networks. This capability facilitates seamless migration from GbE to 10 GbE and allows for gradual or phased upgrades. Now with a substantial and growing ecosystem of cost-competitive switches and servers, 10GBASE-T is poised for rapid growth that is bringing 10 GbE to volume mainstream servers.

10GBASE-T Interconnect Advantages

Data center administrators have a range of 10 GbE interconnect options, including 10GBASE-SR (short-reach fiber optic), enhanced small form factor pluggable (SFP+) direct-attach copper (DAC), and 10GBASE-T. Each of these choices has advantages, but 10GBASE-T's compatibility with existing structured cabling installations and lower-speed equipment make it uniquely suited for broad deployment.

These capabilities combine with superior cost and reach characteristics to make 10GBASE-T the easiest path for migrating from GbE to 10 GbE. The 10GBASE-T standard supports up to 100 meters of reach, which is sufficient for the vast majority of data center deployments, while passive DAC (at a maximum reach of 7 meters (~23 feet)) is best suited to top-of-rack (ToR) implementations. 10GBASE-T copper cabling is also significantly more cost-effective than either fiber or DAC.

With the introduction of 10GBASE-T-enabled equipment such as the Cisco Nexus® 5596T switch, the Cisco Nexus® 2232TM-E 10GE Fabric Extender (FEX), the Cisco Nexus® 7000 1/10GBASE-T line card, and the Intel® Ethernet Controller X540, data center operators can enjoy the many advantages of 10GBASE-T across their infrastructures.

- **Cost-effective twisted-pair copper cabling** for lowest cost 10 GbE deployment
- **Auto-negotiable backward-compatibility** with previous-generation BASE-T networks for a seamless migration to 10GbE
- **Field-terminable twisted-pair cabling** with a familiar RJ-45 interface
- **Existing expertise** based on prior-generation BASE-T knowledge and training
- **Cable lengths up to 100 meters**, sufficient to support nearly all data center topologies (see Figure 1)

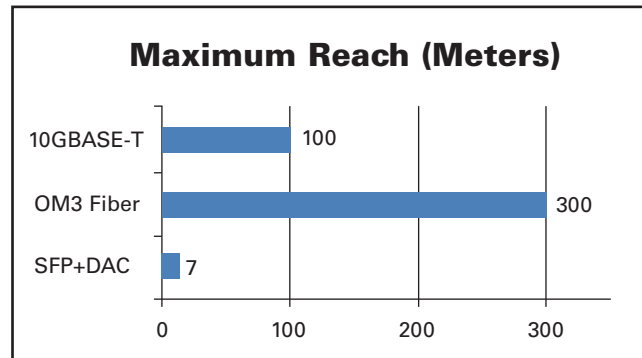


Figure 1. Maximum reach of various 10 Gigabit Ethernet cabling options.

10GBASE-T Reduced Power Requirements

Early 10GBASE-T products were relatively inefficient in terms of power use, with first-generation, single-port adapters consuming nearly the 25-watt maximum allowed by the PCI Express* specification. Design enhancements and advanced production processes, however, have resulted in a new generation of power-efficient 10GBASE-T products.

The Intel Ethernet Controller X540, for example, integrates the media access controller (MAC) and physical layer (PHY) controller into a single package, helping lower power consumption to approximately 6.5 watts per port. Standalone 10GBASE-T PHYs have seen similar improvements in power utilization, resulting in greater port density for 10GBASE-T switches. Many of these PHYs also incorporate a short-reach mode that can reduce power consumption down to 2.5 watts for links of 30 meters (~98 feet) or less (as specified in the IEEE 802.3an standard).

On a watts-per-gigabit basis, 10GBASE-T is significantly more power-efficient than 1000BASE-T, with a 10GBASE-T server adapter port requiring approximately .87 watts versus 1.25 watts for a 1000BASE-T port, an improvement of nearly 70 percent.¹

Further reductions are expected, based on continual advances in manufacturing and the move from 40nm to 28nm semiconductor manufacturing process technology. Future implementation of Energy Efficient Ethernet (EEE, which is already used in GbE adapters) in 10GBASE-T products is expected to further decrease power consumption.

10GBASE-T on the Motherboard

With a low-power, single-chip solution now available, 10GBASE-T has made its way to mainstream server motherboards, and a number of server manufacturers are shipping systems with integrated or modular 10GBASE-T LOM connections based on the Intel Ethernet Controller X540. These connections eliminate the need for add-in PCI Express 10GBASE-T adapters, reducing costs, while backward-compatibility with GbE equipment facilitates upgrades to 10 GbE.

“The introduction of lower-powered, attractively priced 10GBASE-T switches and adapters, in conjunction with the ramp of the more powerful Intel® Xeon® processor E5 family-based servers, should drive very strong adoption of this technology in the coming years.”

– Seamus Crehan, President, Crehan Research

According to multiple research firms, 10GBASE-T will experience significant growth over the next several years. The Dell’Oro forecast shown in Figure 2, for example, indicates that by 2016, 10GBASE-T will be the dominant 10 GbE connection. Major drivers of this growth are 10GBASE-T LOM and the associated switch ecosystem, in addition to products such as 802.11ac wireless access points and high-end workstations.




Figure 2. Projected relative growth of various types of 10 Gigabit Ethernet connectivity.²

10GBASE-T Implementation and Network Equipment 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 technology. Cisco, Intel, and Panduit each deliver products that together form a comprehensive, integrated 10GBASE-T ecosystem solution that has been tested to ensure interoperability, availability, and reliability for high-speed 10 GbE networks. This cost-effective and flexible approach can be developed through the integration of 10GBASE-T enabled devices, including the following:

- Cisco Nexus 2232TM-E 10GbE fabric extender
- Cisco Nexus 5596T switch, part of the Cisco Nexus 5000 Series family of switches
- Cisco Nexus 7000 1/10GBASE-T line card, part of the Cisco Nexus 7000 Series platform
- Intel Ethernet Controller X540 and Intel® Ethernet Converged Network Adapter X540
- Panduit TX6A™ 10Gig™ Copper Cabling System

This solution can be deployed in 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 topology offers distinct advantages for both new and existing data centers migrating from GbE to 10 GbE.

As part of the natural progression of technology improvements, the next generation of 10GBASE-T equipment from Cisco has improved performance and reliability characteristics, which in turn create a solid foundation for multi-protocol traffic environments, including Fibre Channel over Ethernet (FCoE). Intel’s current generation of 10GBASE-T products also supports multiple storage protocols, including FCoE.

Cisco Nexus® 2232TM-E 10GbE Fabric Extender

The Cisco Nexus 2232TM-E Fabric Extender, shown in Figure 3, provides 32 ports of 1000BASE-T/10GBASE-T connectivity and eight SFP+ ports; these devices are designed to behave as remote I/O modules for a parent Cisco Nexus 5500 Series or 7000 Series switch. The newer generation 10GBASE-T PHYs (40nm) used in the Cisco Nexus 2232TM-E bring improved bit error rate (BER) characteristics and reduced power consumption (~210 watts at 30 meters and 240 watts at 100 meters). The fabric extender, which acts as a line-card extension of the Nexus switch fabric, works together with the Nexus parent switch to simplify operational management by enabling all ports to be managed as if they were local interfaces.

The Cisco Nexus 2232TM-E forwards all traffic to the parent switch over 10 GbE SFP+ uplinks. 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 GbE ports and more than 1,000 10 GbE ports. This distributed, modular system architecture provides the following benefits:

- Architectural flexibility
- Highly scalable server access (GbE/10 GbE ready anytime)
- Simplified operations and management
- Increased business benefits, including fabric-extender transceivers for cost-effective uplink to fabric




Figure 3. Cisco Nexus® 2232TM-E fabric extender.

Cisco Nexus® 5596T Switch

The Nexus 5596T, shown in Figure 4, is the newest addition to the Nexus 5500 Series family and the first one to be compatible with 10GBASE-T, meeting the need associated with the growing prevalence of 10GBASE-T LOM connections being integrated into mainstream server boards for high-performance, high-density switching infrastructure. Part of Cisco's Unified Fabric, the Nexus 5596T offers the opportunity to take advantage of the simplicity, speed, and design flexibility of 10GBASE-T technology. The Nexus 5596T switch is based on the same application-specific integrated circuit (ASIC) and chassis as the Nexus 5596UP switch, but it provides 32 fixed 1000BASE-T/10GBASE-T ports, and through the addition of the new 12-port 10GBASE-T Cisco Generic Expansion Modules (GEMs), it can support up to 68 total 10GBASE-T ports in a two-RU (rack unit) design.

Moreover, customers can deploy any of the GEMs in any of the Nexus 5596T's three GEM slots. The Nexus 5596T also includes 16 fixed 10 GbE SFP+ capable ports, which can be used to connect to aggregation switches, servers, or Nexus 2000 Fabric Extenders such as the Nexus 2232TM-E. The switch supports 10 GbE (fiber and copper), Fibre Channel, and FCoE, offering up to 1,920 gigabits per second (Gbps) of throughput and up to 96 ports. In addition to three expansion modules, the Cisco Nexus 5596T supports a Layer 3 module that provides up to 160 Gbps of Layer 3 forwarding capability (240 million packets per second) that can be shared by all the I/O ports in the chassis. The 5596T also supports front-to-back and back-to-front airflow.



Figure 4. Cisco Nexus® 5596T switch.

In conjunction with the Cisco Nexus 2232TM-E Fabric Extender, the Cisco Nexus 5596T can be used as a high-density GbE/10 GbE switching system, converging hundreds of 10 GbE connections into a single management plane. The highly scalable, distributed, modular system scales up to 768 ports of 1000BASE-T/10GBASE-T connectivity (24 fabric extenders are currently supported per Cisco Nexus 5500 platform). Combining low-latency cut-through switching, low power usage, and support for PortChannel and virtual PortChannel (vPC) for server access, the Cisco Nexus 5596T and Nexus 2232TM-E Fabric Extender provide outstanding connectivity options for the bare-metal and virtualized server edge.

Cisco Nexus® 7000 1000BASE-T/10GBASE-T Line Card

The Nexus 7000's enhanced F2-Series 1000BASE-T/10GBASE-T line card is a 48-port module that offers outstanding flexibility and wire-rate performance on each port. It is a high-performance and high-density 10GBASE-T Ethernet module designed for mission-critical data center networks. The Cisco Nexus 7000 F2e Copper Module provides the highest density of wire-rate 10GBASE-T Ethernet ports on the market,3 with up to 768 wire-rate 10GBASE-T Ethernet ports directly attached to a single Cisco Nexus 7000 18-slot chassis, as shown in Figure 5.

The Cisco Nexus 7000 F2e Copper Module offers customers low 10GBASE-T power consumption, with less than 9 watts per port for typical deployment scenarios. This power efficiency is enabled in part by the module's switch-on-chip (SoC) architecture, in which a single ASIC implements all the module functions. Each SoC manages four front-panel interfaces. In addition, power consumption per link can be further improved by enabling the Energy-Efficient Ethernet feature.



Figure 5. Cisco Nexus® 7018 chassis and 48-port 1000BASE-T/10GBASE-T module.

Further benefits of the Nexus 7000 F2e copper line card include FCoE-readiness, IEEE 802.1ae MACsec encryption on all 48-ports, GbE/10 GbE auto-negotiation on each port (enabling a seamless upgrade to 10 GbE), and reduced deployment costs with passive cabling and RJ-45 connectors, along with all the features of the Nexus 7000 modular platform and NX-OS operating system.

Intel® Ethernet Controller X540 and Intel® Ethernet Converged Network Adapter X540

The Intel Ethernet Controller X540, shown in Figure 6, is the industry's first single-chip 10GBASE-T controller, combining MAC and PHY into a single package. This controller is designed for low-cost, low-power LOM and converged network adapter solutions.



Figure 6. Intel® Ethernet Controller X540.

Key features include the following:

- **Virtual Machine Device Queues (VMDq) and support for PCI-SIG Single-Root I/O Virtualization (SR-IOV)** provide optimized I/O virtualization for virtualized and non-virtualized servers
- **Unified networking** delivers LAN, NFS, iSCSI, and FCoE in a single, low-cost adapter, eliminating the need for dedicated storage adapters
- **Power-efficient 10GBASE-T performance** offers advantages in cost, power, and server-slot savings compared to the use of multiple GbE adapters
- **Backward-compatibility with 1000BASE-T infrastructures** helps ensure easy deployment and upgrading
- **Dual-port design** for redundancy and greater throughput
- **Reliable, consistent performance** across a broad range of operating systems and hypervisors

The Intel® Ethernet Converged Network Adapter X540-T2 and T1 are powered by the Intel Ethernet Controller X540 and are Intel's fourth-generation 10GBASE-T adapters, as shown in Figure 7.



Figure 7. Intel® Ethernet Converged Network Adapter X540-T2.

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

The development of 10GBASE-T equipment has been accompanied by advances in cabling technology. Early generations of Category 6A 10 GbE-compliant cabling were relatively large in diameter (as much as 0.35 inches), to suppress coupling between neighboring cables (also 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 0.30 inches without sacrificing performance. For high-density applications such as large data centers in which reduced cabling density is important, Panduit also 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 feet) with a cable diameter of only 0.61 centimeters (0.24 inches). This size is comparable to that of previous

generations of Category 6 cabling and allows full reuse of existing pathways and cable managers to further facilitate seamless migration from GbE to 10 GbE. A space utilization comparison of 100 cables of standard Category 6A, Panduit TX6A, and Panduit TX6A-SD is shown in Figure 8.




Figure 8. Comparative representation of space-utilization by 100 cables of various types.

10GBASE-T Network Architecture

10GBASE-T brings architectural flexibility that can support topologies commonly used in data centers, including ToR as well as the more traditional EoR and MoR. Each of these architectures has distinct benefits and constraints, so the choice among them for a particular implementation is typically based on the preferences of the data center manager, the existing facilities, equipment layout, and the applications being run. The following sections present the advantages of each architecture and how each can be used in the efficient implementation of 10GBASE-T.

Top-of-Rack (ToR) Network Topology

ToR is a distributed modular approach, shown in Figure 9, in which a switch or a fabric extender such as the Cisco Nexus 2232TM-E resides at the top of each server rack. In this architecture, 10GBASE-T provides the flexibility of connecting existing 1000BASE-T or older switches and servers to newer 10 GbE equipment using inexpensive Category 6A patch cords. SFP+ DAC, by contrast, requires more expensive cables that are not compatible with 1000BASE-T switches.

As the number of cables exiting the rack increases with the number of servers, cost and space constraints become more pressing. When used in a ToR architecture, the Nexus 2232TM-E fabric extender aggregates network connections in the rack, reducing the number of intra-enclosure cable runs back to the Cisco Nexus 5000 or 7000 Series switches. Each Nexus 2232TM-E has 32 downlink 1000BASE-T/10GBASE-T ports, allowing it to support up to 32 servers. A common configuration includes 12 cabinets per row, with each cabinet containing up to 16 2U servers or 32 1U servers, and each server equipped with dual-port Intel Ethernet Converged Network Adapters X540-T2 or LOM connections. One or two Nexus 2232TM-E fabric extenders are placed at the top of each cabinet. The second fabric extender can be used either to provide redundancy or to enable more servers to be added per cabinet.




Figure 9. Top-of-rack (ToR) network topology.

In this scenario, the Nexus 2232TM-E downlinks to the Intel Ethernet Converged Network Adapters X540-T2 in the server cabinet through Panduit TX6A™ 10Gig™ or TX6A-SD™ 10Gig™ Patch Cords. The Nexus 2232TM-E fabric extender acts as a remote line card for a Cisco Nexus 5000 or 7000 Series switch, which is located in a switch cabinet either at the end of the row or in the middle of the row. For redundancy, two Nexus 5000 or 7000 Series switches are used, generally deployed in separate cabinets. The Nexus 2232TM-E is uplinked to the aggregation switch using Panduit OM3 structured fiber or SFP+ DAC, depending on the location of and distance to the switch. Locating the Nexus 5000 or 7000 in the middle of the row reduces cable lengths to the Nexus 2232TM-E and allows the use of SFP+ DAC cabling. 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.

Business Needs and Situations Well-Suited to ToR Topology with 10GBASE-T:

- New data centers without any existing cabling or equipment
- Need for scalability and ease of adding cabinets
- Ease of migration from 10 Gbps networking toward 40Gbps
- High-density layouts using large numbers of 1U servers per cabinet
- Reduced cable congestion
- Modular, pre-configured cabinets

There are eight SFP+ uplink ports on the Nexus 2232TM-E, for a 4:1 oversubscription ratio. To provide greater redundancy, four uplinks can be connected to one aggregation switch, with the remaining four uplinks connecting to a different switch located either in the same cabinet or a different cabinet. One Nexus 5000 Series switch can support up to 24 Nexus 2232TM-E fabric extenders, but best practices call for that number to be reduced to 12, allowing for redundancy. Therefore, a typical row with 10 to 12 server cabinets, each with 32 1U servers and two Nexus 2232TM-E fabric extenders, would be supported by a switch cabinet with two Nexus 5000 Series switches. Similarly, a row of 18 server cabinets each with 32 1U servers and two Nexus 2232TM-E fabric extenders would have a switch cabinet with four Nexus 5000 Series switches or two switch cabinets each with two Nexus 5000 Series switch. Alternately, one Cisco Nexus 7000 chassis can support up to 48 fabric extenders.

End-of-Row (EoR) Network Topology

Standard structured cabling practices are used in the EoR topology to provide efficient cable organization, allowing the user to take optimal advantage of 10GBASE-T. In this topology, inexpensive Category 6A cabling can be used for the vast majority of connections between switches, fabric extenders, and servers, due to the 100-meter reach and backward compatibility of 10GBASE-T. The use of RJ-45-based patch panels also simplifies management and changes to the environment, while helping to reduce the potential for airflow blockage.

In EoR topology, a Nexus 7000 chassis with 1000BASE-T/10GBASE-T modules or a Nexus 5596T Series switch is located in a switch cabinet at either end of the row. In this configuration, servers are directly connected at wire-rate to the switch ports through a structured cabling infrastructure that may include passive patch panels at the top of each server rack. The oversubscription ratio can be managed in the modular chassis by simply adding or removing fiber-uplink line cards.

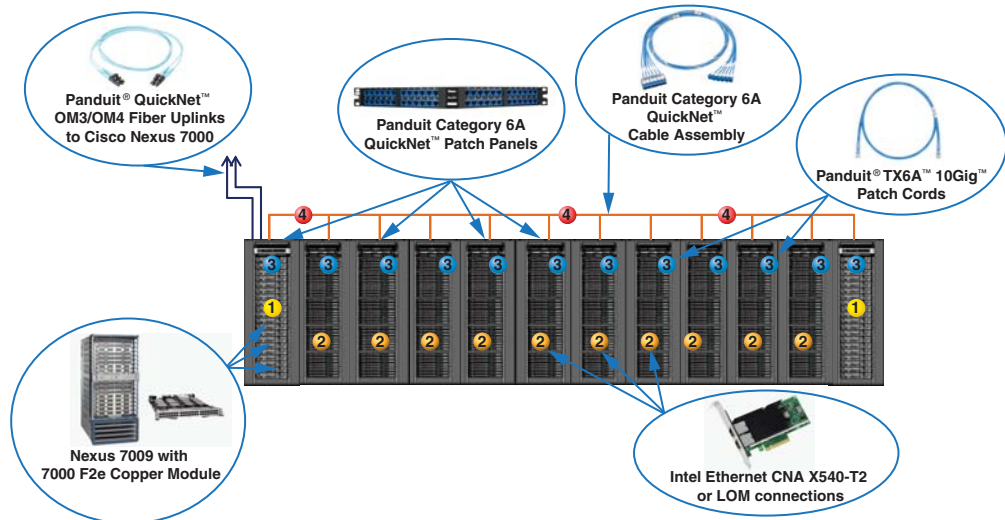
Business Needs and Situations Well-Suited to EoR Topology with 10GBASE-T:

- 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

For example, a row with 12 server cabinets each hosting 32 servers would connect to eight Nexus 7000 F2e copper modules in an 18-slot chassis placed at either end of the row. The chassis would then contain another eight fiber line cards for wire-rate uplinks, four fiber line cards enabling 2:1 oversubscribed uplinks, or two fiber line cards enabling 4:1 oversubscribed uplinks to the data center core or a campus edge router.

A practical deployment of the EoR model with the Nexus 7000 F2e Copper Module is the migration of an existing GbE topology. In this scenario, the data center can retain its structured cabling and replace its GbE-capable EoR chassis with a comparably-sized Nexus 7009 chassis. This approach allows the replacement of the aggregation switch and takes advantage of the dual-rate GbE/10 GbE capability of the F2e Copper Module to progressively upgrade server equipment with the latest 10GBASE-T converged network adapter or LOM offerings.

The EoR switch cabinet or a cabinet adjacent to the switch cabinet generally contains multiple patch panels, called “patch fields,” for connection to all of the server cabinets. As previously mentioned, the server cabinets in the EoR architecture, which is shown in Figure 10, 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 Converged Network Adapters X540-T2 or LOM connections 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 small diameter TX6A-SD™ 10Gig™ UTP Copper Cable. To minimize congestion, the Panduit Net-Access™ CabRunner® Overhead Cable Routing System can be used with horizontal cable managers.



- 1 Nexus 7009 with 7000 F2e Copper Modules located in the switch cabinets located at the end of the row
- 2 Intel Ethernet CNA X540-T2 or LOM connections located in each server
- 3 Panduit® TX6A™ or TX6A-SD™ 10Gig Patch Cords connect Intel Ethernet CNA X540-T2 or LOM connections to Panduit Patch Panels at the top of each server cabinet and Nexus 7009 to Patch Panels at the top of the switch cabinets
- 4 Panduit® TX6A™ or TX6A-SD™ 10Gig Cable and Jack Modules or QuickNet™ Pre-terminated Cable Assemblies connect the Panduit Patch Panels in the server cabinets to the Patch Panels in the switch cabinets

Figure 10. End-of-row (EoR) network topology.

Middle-of-Row (MoR) Network Topology

MoR topology is similar to EoR architecture, except that the switch cabinets are located in the middle of the row, with cabling branching to each of the server cabinets, as shown in Figure 11. The cable lengths are generally shorter, which can reduce cabling complexity and allow more connectivity options. The shorter runs can often use lower-cost copper instead of fiber and allow for single patch fields instead of the dual-point patching normally used in EoR. This characteristic of MoR topologies make them relatively simple to maintain. At the same time, however, centralized switching cabinets necessitate proper thermal planning during the layout phase to prevent power and cooling issues.

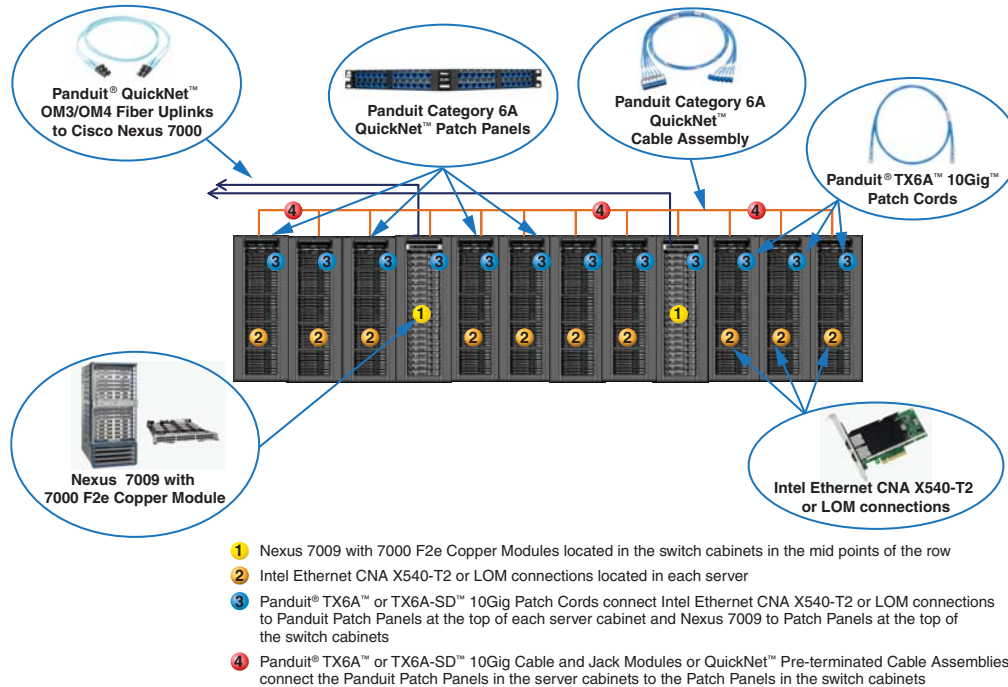


Figure 11. Middle-of-row (MoR) network topology.

The MoR topology is well-suited to pod architectures like the one illustrated in Figure 12, which provides good redundancy and compartmentalized or segregated computing, as commonly required by co-location sites and larger data centers.




Figure 12. Top view of a typical end-of-row (EoR) or middle-of-row (MoR) layout showing redundant cabling.

As with EoR, MoR topologies provide ease of replication and allow simple migration from legacy modular server-access architectures. These topologies also allow for central patching, which reduces oversubscription and optimizes port utilization, an effective potential means of reducing capital expenditures (CAPEX). In addition, because all the switches are in just one or two cabinets, EoR and MoR simplify network management and consolidate security into fewer locations.

Benefits of 10GBASE-T Structured Cabling

Most data centers use structured cabling because it provides an efficient, cost-effective means for managing cabling while providing multiple rerouting points to handle changes or replacement in case of equipment failure. In the EoR and MoR architectures, Panduit TX6A™ 10Gig™ Patch Cords connect the Intel Ethernet Converged Network 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™ Patch Panels populated with pre-terminated cable assemblies or DP6A™ 10Gig™ Punchdown Patch Panels, come with 24, 32, or 48 10 GbE ports (chosen 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 cables terminated with Mini-Com® TX6A™ 10Gig™ Jack Modules on both ends connect the patch panels in the server cabinet to the patch field in the switch cabinet. This cable assembly is called horizontal cabling or “permanent link” and has a maximum reach of 90 meters (~295 feet) in a 10GBASE-T network. The maximum 10GBASE-T channel length of 100 meters is obtained by adding up to 10 meters (~33 feet) of TX6A™ 10Gig™ Patch Cords. The horizontal cable can be terminated during installation if precise lengths are needed, or they can be purchased as pre-terminated cable assemblies, such as Panduit’s QuickNet™ Cabling System. The patch panel in the switch cabinet is then connected to the Cisco Nexus 2232TM-E fabric extender through another Category 6A patch cord.

Conclusion

10GBASE-T, which can be used with all standard data center topologies (including ToR, EoR, and MoR), is well-suited for widespread use in today's networks. It simplifies data center infrastructure by providing high flexibility for new data centers and the easiest upgrade path for existing ones. By offering high reach and bandwidth at low cost, 10GBASE-T meets the demands of virtualized and ever-flattening networks. Building on its proven BASE-T legacy and new LOM implementations, 10GBASE-T now offers the lowest cost of deployment for 10 GbE, with expected continuing decreases in power consumption. 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 provides products that transport data, voice, and video within buildings and around the world. As a worldwide leader in networking, Cisco is dedicated to transforming how people connect, communicate, and collaborate. Information about Cisco can be found at www.cisco.com. For ongoing news, please visit newsroom.cisco.com.



170 W. Tasman Dr.
San Jose, CA 95134
(800) 553-6387



2200 Mission College Blvd.
Santa Clara, CA 95054
(408) 765-8080

About Intel

Intel develops technologies, products, and initiatives to continually advance how people work and live. Intel is the industry leader in Ethernet adapters and has delivered quality Ethernet products for over three decades.

Learn more at www.intel.com/go/ethernet.

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

ZCAN03–WW-ENG 7/2013

Copyright and Trademark Information

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

MTP is a registered trademark of US Conec Ltd.

NOTE: This is an application note for Panduit Infrastructure and is not a requirement to install or operate the Cisco Nexus 7000 Series switches or I/O modules.



18900 Panduit Drive
Tinley Park, IL 60487
(800) 777-3300

- 1 Figures based on the latest Intel® Ethernet 1000BASE-T Adapter and the Intel® Ethernet 10GBASE-T Adapter as of June 2013.
- 2 Source: Dell'Oro Group. Ethernet Switch Five-Year Forecast, January 2013.
- 3 As of June, 2013.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1005R)

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.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site <http://www.intel.com/>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.

*Other names and brands may be claimed as the property of others.

Copyright© 2013 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

0613/BY/MESH/PDF

329050-001US