



## WHITE PAPER

# 10 GIGABIT ETHERNET SWITCHING FOR ENTERPRISES TECHNOLOGY OVERVIEW AND APPLICATIONS

Important technology advances and significant price and performance improvements have enabled 10 Gigabit Ethernet to be deployed not only in data centers but also throughout campus networks. Broader deployments of 10 Gigabit Ethernet are being accelerated by increasing bandwidth requirements and the aggregate growth of enterprise applications, examples of which are discussed in this paper.

## INTRODUCTION

Since the IEEE 802.3ae standard was ratified in mid-2002, 10 Gigabit Ethernet port shipments have grown from hundreds of ports per quarter to tens-of-thousands of ports per quarter. This rapid growth in 10 Gigabit Ethernet deployments can be attributed to a number of factors, including:

- **Significant 10 Gigabit Ethernet Price-per-Port Improvements**—Current 10 Gigabit Ethernet pricing is now less than one-fifth the pricing in mid-2002. As a result, 10 Gigabit Ethernet price and performance today, including cost of optics, is comparable to Gigabit Ethernet-over-fiber price and performance in intelligent modular switches.
- **New Optics have Enabled Broader 10 Gigabit Ethernet Deployments**—The availability of new optics now enables 10 Gigabit Ethernet to be deployed anywhere from the data center to the wiring closet, using existing fiber cabling.
- **Increasing Bandwidth Factors**—First, Gigabit Ethernet-to-desktops deployments have grown to several million ports per quarter by the end of 2004. This broad adoption has significantly increased the oversubscription ratios of the rest of the network. 10 Gigabit Ethernet can help bring these oversubscription ratios back in line with network-design best practices. Second, server adapter and PCI bus advancements have enabled servers to generate more than 7 Gbps of traffic, increasing demand for 10 Gigabit Ethernet connectivity to servers. Finally, new applications are accelerating the need for 10 Gigabit Ethernet performance throughout the campus, within a data center, and between data centers. These applications are described in more detail in later sections.

These factors are expected to continue to fuel the momentum of the 10 Gigabit Ethernet market, which is expected to rapidly grow from US \$385 million in 2004 to US\$2.9 billion in 2009, according to the Dell'Oro Group.

This paper discusses the 10 Gigabit Ethernet technologies and applications that are relevant for enterprise networks. 10 Gigabit Ethernet technologies and applications for service provider deployments (such as WAN PHY and intra-point of presence [PoP] interconnects) are beyond the scope of this paper.

## TECHNOLOGY OVERVIEW

### MAC-Layer Attributes

Because 10 Gigabit Ethernet is Ethernet, it takes advantage of the wealth of Ethernet technologies that have been developed over the years and simplifies the migration to this higher-speed technology. Just like Fast Ethernet and Gigabit Ethernet before it, 10 Gigabit Ethernet uses the IEEE 802.3 Ethernet MAC protocol, Ethernet frame format, and frame size. It supports standard Ethernet services, such as 802.3ad link aggregation, enabling up to eight 10 Gigabit Ethernet links to be aggregated into a virtual 80-Gbps connection. Because 10 Gigabit Ethernet is also full duplex point-to-point technology, it can support simultaneous traffic from both ends of the link without any packet collisions. Therefore it does not have inherent distance limitations. The maximum link distances are determined by the physics of transmission and the physical media optics and not by the diameter of an Ethernet collision domain.

## Physical-Layer Attributes

### General Interface Naming Convention and Operating Ranges

One of the first questions asked when a new Ethernet technology is introduced is “how far can it go?”. The answer for 10 Gigabit Ethernet, like the Ethernet technologies before it, depends on the physical interface type used. The wide range of available 10 Gigabit Ethernet physical interface options requires a general naming convention to make sense of the various optics, fiber types, and distances (Table 1).

10 Gigabit Ethernet physical-layer interfaces tend to use the following general naming convention:

- **Prefix** = “10GBASE-” = 10-Gbps baseband communications
- **First suffix** = Media type or wavelength, if media type is fiber
- **Second suffix** = PHY encoding type
- **Third suffix** = Number of Wide Wave Division Multiplexing (WWDM) wavelengths or XAUI lanes

**Table 1.** General Physical Interface Naming Convention for 10 Gigabit Ethernet

Prefix	First Suffix = Media Type or Wavelength	Second Suffix = PHY Encoding Type	Third Suffix = Number of WWDM Wavelengths or XAUI Lanes
<b>10GBASE-</b>	Examples: <b>C</b> = Copper (twin axial) <b>S</b> = Short (850 nm) <b>L</b> = Long (1310 nm) <b>E</b> = Extended (1550 nm) <b>Z</b> = Ultra extended (1550 nm)	Examples: <b>R</b> = LAN PHY (64B/66B) <b>X</b> = LAN PHY (8B/10B) <b>W</b> = WAN PHY (64B/66B)	Examples: If omitted, value = 1 (serial) <b>4</b> = 4 WWDM wavelengths or 4 XAUI lanes

For example, a 10GBASE-LX4 optical module uses a 1310-nanometer (nm) laser, LAN PHY (8B/10B) encoding, and four WWDM wavelengths. A 10GBASE-SR optical module uses a serial 850-nm laser, LAN PHY (64B/66B) encoding, and one wavelength. The IEEE 802.3an task force is targeting to finalize the standard for 10 Gigabit Ethernet over twisted pair copper cabling (10GBASE-T) by late 2006.

Table 2 summarizes the operating ranges and media types supported for various 10 Gigabit Ethernet interfaces that would be used in enterprise deployments.

**Table 2.** 10 Gigabit Ethernet Operating Ranges

10 GE Physical Interface	Typical Deployment	Operating Range Over:			
		62.5 micron Multi-Mode Fiber (FDDI-grade)	50 micron Multi-Mode Fiber (MMF)	10 micron Single Mode Fiber (SMF)	Twinaxial Copper
<b>10GBASE-CX4</b>	Data Center	–	–	–	15m
<b>10GBASE-SR</b>	Data Center	26m–33m	66m–300m	–	–
<b>10GBASE-LX4</b>	Campus Or Data Center	300m	240m–300m	–	–
<b>10GBASE-LR</b>	Campus Or Metro	–	–	10 km	–
<b>10GBASE-ER</b>	Metro	–	–	40 km	–
<b>10GBASE-ZR</b>	Metro Or Long-Haul	–	–	80 km	–
<b>DWDM</b>	Metro Or Long-Haul	–	–	80 km—32 wavelengths over single-strand SMF	–

More than 75 percent of existing fiber cabling from campus distribution layer to the wiring closet is FDDI-grade (62.5 micron) multimode fiber (MMF) and the distance requirements are typically greater than 100 meters (m). Thus, deploying 10 Gigabit Ethernet to wiring closets over existing FDDI-grade MMF will typically require the 10GBASE-LX4 optic.

## Form Factors

10 Gigabit Ethernet pluggable interfaces are available in a number of form factors, such as XENPAK, X2, and XFP. From a deployment standpoint, the primary differences among these options are 1) the breadth of 10 Gigabit Ethernet physical interfaces supported in a given form factor and 2) physical size. For example, the XFP form factor currently does not support 10GBASE-LX4 nor 10GBASE-CX4 optics because of space constraints. The various form factor options are optically interoperable with each other so long as the 10 Gigabit Ethernet physical interface type (for example, 10GBASE-LX4 or 10GBASE-SR) is the same on both ends of the link.

## 10 Gigabit Ethernet Advantages vs. Aggregating Multiple Gigabit Ethernet Links

Many network managers are weighing the option of using Gigabit Ethernet link aggregation as opposed to deploying a single, 10 Gigabit Ethernet link. As always, there are tradeoffs associated with each option. However, 10 Gigabit Ethernet provides some important advantages over aggregating multiple Gigabit Ethernet links:

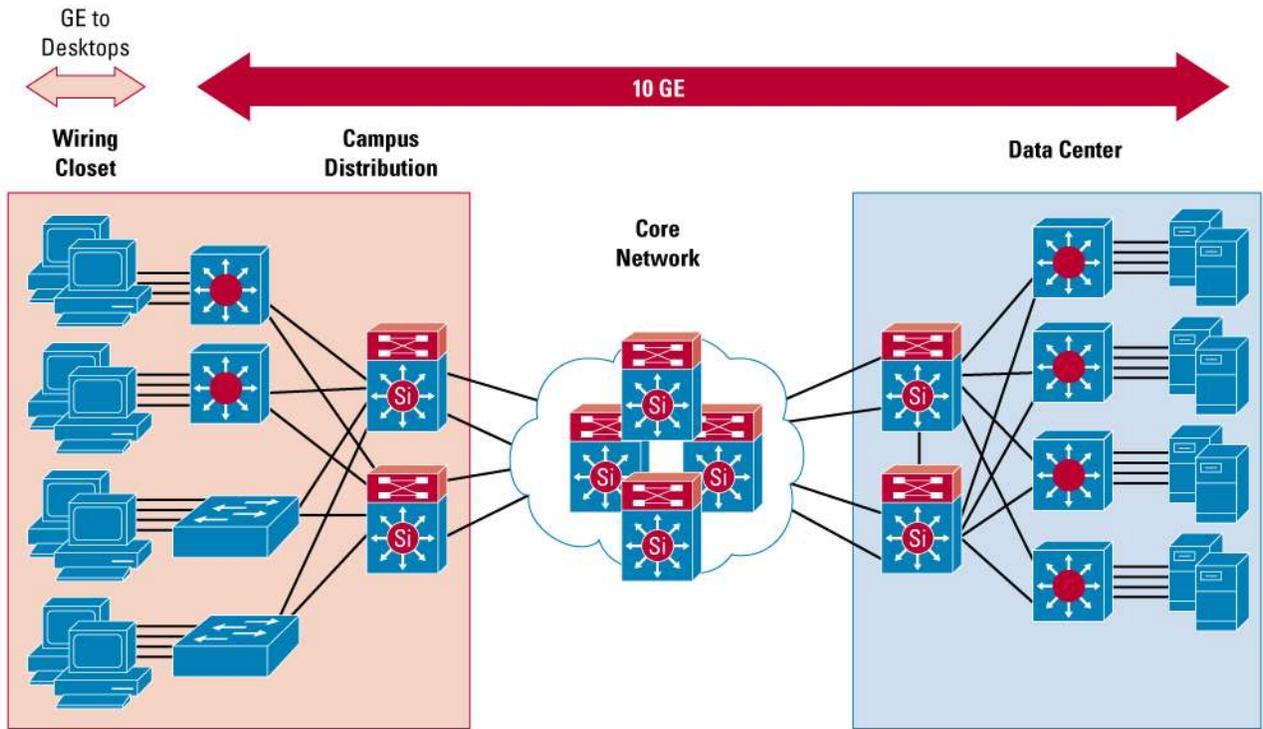
- **Less Fiber Usage**—A 10 Gigabit Ethernet link uses fewer fiber strands compared with Gigabit Ethernet aggregation, which uses one fiber strand per Gigabit Ethernet link. This 10 Gigabit Ethernet advantage reduces cabling complexity in data centers and more efficiently uses existing fiber cabling in campus environments where laying additional fiber could be cost-prohibitive.
- **Greater Support for Large Streams**—Traffic over aggregated 1 Gigabit Ethernet links can be limited to 1 Gbps streams because of packet sequencing requirements on end devices. 10 Gigabit Ethernet can more effectively support applications that generate multigigabit streams due to the greater capacity in a single 10 Gigabit Ethernet link.
- **Longer Deployment Lifetimes**—10 Gigabit Ethernet provides greater scalability than multiple Gigabit Ethernet links, enabling longer deployment lifetimes. Up to eight 10 Gigabit Ethernet links can be aggregated into a virtual 80-Gbps connection.

## 10 GIGABIT ETHERNET ENTERPRISE APPLICATION SCENARIOS

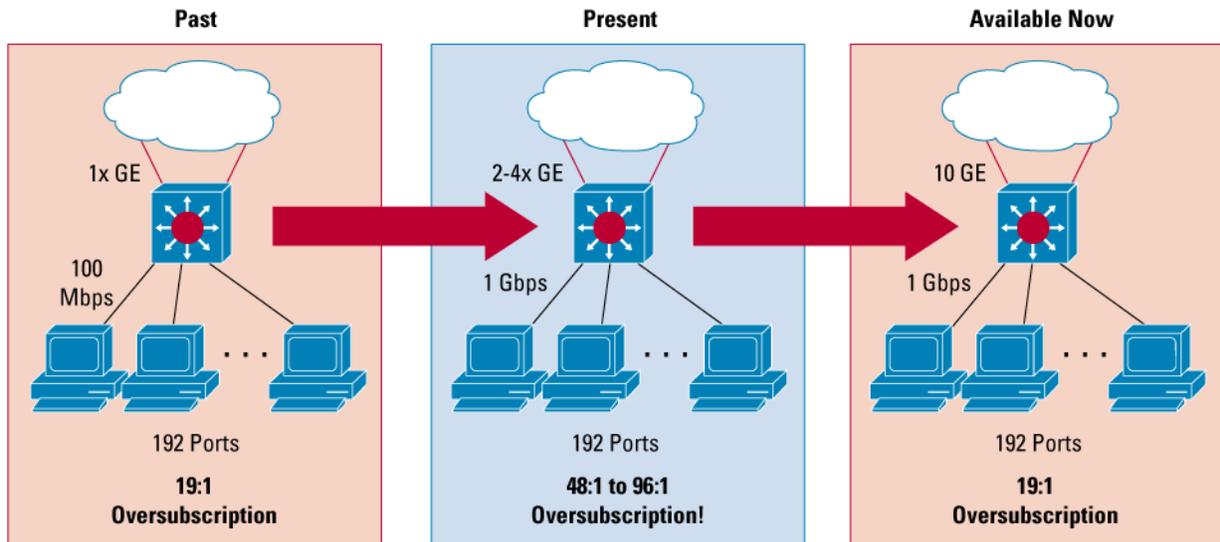
As previously mentioned, 10 Gigabit Ethernet can now be deployed over existing fiber cabling from the data center to the wiring closet uplinks (see Figure 1). 10 Gigabit Ethernet deployments continue to extend beyond the network core to improve network scalability as end devices increase their bandwidth connectivity. For example, Gigabit Ethernet-to-desktops deployments have grown to several million ports per quarter by the end of 2004. This broad adoption has significantly increased the oversubscription ratios of wiring closet uplinks, especially because more than 90 percent of wiring closet traffic flows north to south through the uplinks.

Figure 2 shows the evolution of an example high-density campus wiring closet. In the late 1990s, it was common to deploy 10/100 Ethernet to desktops with redundant Gigabit Ethernet uplinks. If there were 192 users per switch, then the oversubscription ratio was roughly 19:1, which is within standard network design best practices of 15:1 to 20:1 wiring closet bandwidth oversubscription. However, as Gigabit Ethernet to desktops has rolled out over the years, these oversubscription ratios have ballooned to 48:1 or 96:1 even when the wiring closet uplinks have been increased to two or four Gigabit Ethernet channels. Deploying 10 Gigabit Ethernet uplinks with today's switching solutions can help bring the wiring closet oversubscription ratios back in line with network design best practices and scale bandwidth capacity for future requirements.

**Figure 1.** 10 Gigabit Ethernet Deployment Throughout the Enterprise



**Figure 2.** Scaling Wiring Closet Uplinks with 10 Gigabit Ethernet

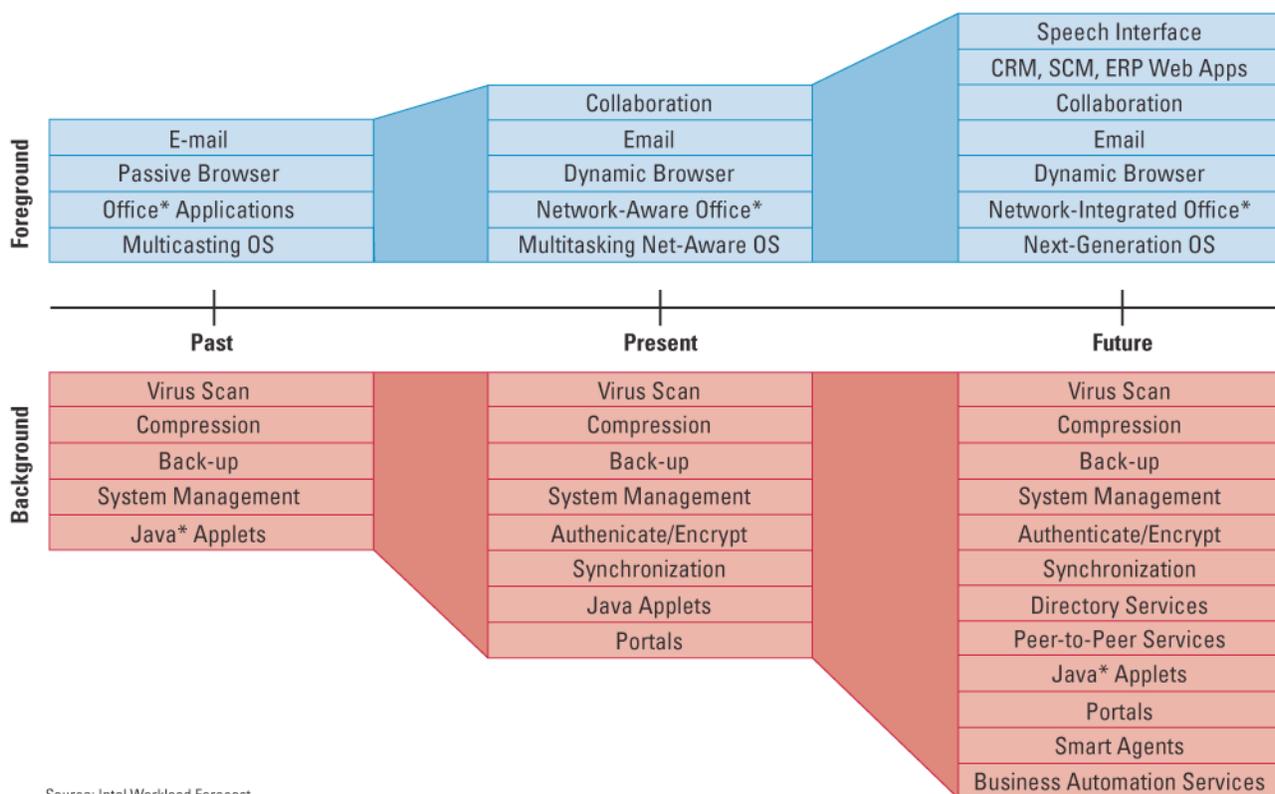


## Desktop Applications

Enterprise-wide 10 Gigabit Ethernet deployments support the continued growth in desktop applications which, in aggregate, is accelerating higher-bandwidth requirements. Examples include:

- Aggregate Desktop Data Workloads**—The aggregate bandwidth consumption per desktop is increasing because of increasing desktop workloads (see Figure 3) and the greater bandwidth requirements of new applications. For example, PC backup applications are critical, especially with rising employee reliance upon recent PC data. Data loss decreases and backup frequency increases when backups are automated instead of user-initiated. Frequent PC backups across all desktops in an organization places continual load on the network especially as file sizes continually increase (for example, Microsoft Outlook data files and PowerPoint presentations). In addition, companies are transitioning from traditional client/server applications (fat, proprietary client on each desktop) to Web-based applications (thin, standard browser on each desktop) to capture the operational and development cost savings associated with Web technologies. However, this transition can result in higher bandwidth usage because browsers may rely more on communicating with servers for intelligence and processing than proprietary clients.

**Figure 3.** Increasing Desktop Workloads



- IP Video Applications**—Enterprises are deploying bandwidth-rich IP video applications to improve productivity and operational costs. For example, e-learning increases employee productivity by providing low-cost, 24-hour access to critical training information, enabling “just-in-time” sales training, quick refreshers on how to deliver a service, lectures, and skills and regulation training. Corporate and executive IP video communications increase corporate alignment to business objectives and strengthen employee morale, and are an especially effective way to increase communication within a global company. IP video surveillance solutions are being deployed to increase security visibility and to accelerate the retrieval and analysis of archived events. IP video conferencing enables efficient collaboration among employees who need to communicate visually but do not have the time to commute to a designated location. Each of these IP video applications can generate numerous multiple-megabit IP video streams, depending on desired video quality, resulting in significant network-bandwidth consumption.

- **Industry-Specific Applications**—Many industries have custom applications that require significant bandwidth capacity and high performance. Whether the application is clustered or based on a client-server model, 10 Gigabit Ethernet can rapidly increase the performance of the network. In the healthcare industry, for example, digital imaging applications (such as Picture Archive Systems [PACS]) are often used to lower the costs and reduce the delay of retrieving and analyzing medical images (such as X-rays, MRIs, and CAT scans), increasing physician and staff productivity. In the media and advertising industries, digital video applications enable companies to efficiently develop video segments and then edit and review them among distributed teams. In the manufacturing industry, large CAD and CAM design files are increasingly being shared among teams located in different locations. And in the financial industry, the continual need for more powerful, real-time financial information continues to elevate network performance requirements.

The aggregate growth of these example applications and other desktop applications is accelerating the need for 10 Gigabit Ethernet performance across the enterprise network.

### Storage Networking

The continuous increase in demand for storage capacity is propelled by applications such as customer care, messaging, e-commerce, rich online media, and catalog content. This information explosion is challenging IT managers to find cost-effective ways to access, manage, and protect this data.

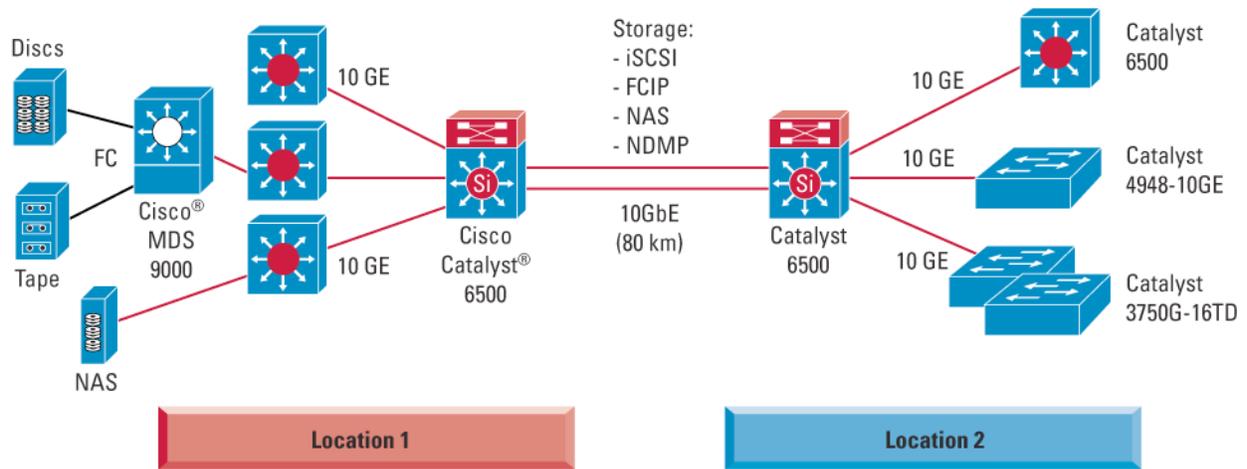
Migrating from server-centric, direct-attached storage to network-centric, shared storage is an important strategy for achieving these goals. The ability to share networked storage in the data center, across the metropolitan area, and across the enterprise provides the following benefits:

- Scaled, shared, and maximized usage of storage and information resources
- Simplified administration of the storage environment
- Minimized total cost of ownership (TCO) for storage
- Improved data availability and integrity

Utilizing 10 Gigabit Ethernet, IT managers can now take their networked storage environments to the next level and use Ethernet-based networking for the most demanding storage solutions, such as:

- **Data Center Backup and Disaster Recovery for Greater Business Resiliency**—Enterprises have been challenged to develop business-continuation and disaster-recovery strategies that are cost-effective, secure, and scalable enough to meet their demanding requirements. An important factor of the move to metropolitan storage networks is the need to establish backups and remote mirrors at remote locations to provide business-continuation and disaster-recovery support for critical data. In addition, companies are also faced with the need to expand data centers that have reached their capacity or alternatively the requirement to centralize data center resources of multiple campuses or locations. The distance capabilities of 10 Gigabit Ethernet allow enterprises to provide high-speed connectivity between locations that are 80 km apart. Distances can be even further extended with the use of optical amplifiers and dispersion compensators. Enterprises can therefore support multiple campuses within this radius, supporting storage-to-server and storage-to-storage data transfers. With the high bandwidth, low latency, and security offered by 10 Gigabit Ethernet and Intelligent Switching, it becomes easier to move data seamlessly between geographically dispersed components of an enterprise storage system. Figure 4 shows a 10 Gigabit Ethernet infrastructure that supports all IP storage-based metro solutions and technologies including Network Attached Storage (NAS), Internet Small Computer System Interface (iSCSI), Fibre Channel over IP (FCIP), and Network Data Management Protocol (NDMP).

**Figure 4.** 10 Gigabit Ethernet for Backup and Disaster Recovery

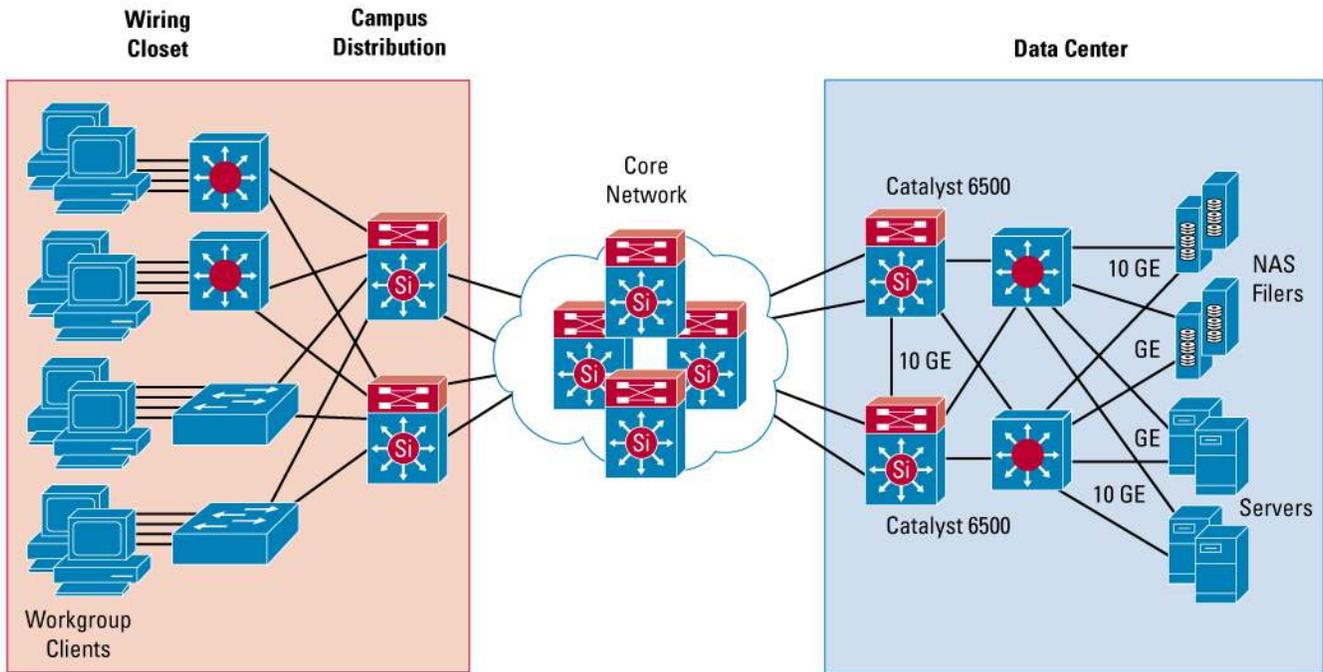


For deployments that require higher bandwidth aggregation, longer distances, low latency, and support for non-IP technologies (such as Fibre Channel or IBM's Enterprise Systems Connection [ESCON] protocol), Dense Wave Division Multiplexing (DWDM) provides high-capacity, protocol-independent access and transport of storage traffic across metropolitan-area network (MANs). Critical storage applications for such optical MAN connectivity include backup, remote mirroring, disaster recovery, clustering, and storage outsourcing. Synchronous mirroring requires very low latency and high bandwidth and 10 Gigabit Ethernet provides the ideal combination of these factors to enable such mission-critical business requirements.

- **Network Attached Storage (NAS) for High-Performance Data Sharing and Storage Consolidation**—NAS has led the way for the mainstream deployment of IP-based storage consolidation and file sharing. NAS has achieved popularity in many environments including collaborative workgroup development, engineering, e-mail, Web serving, and general file serving. Because of the customized nature of their operating systems, NAS filers have been tuned to carry out I/O extremely efficiently so they can easily fill multiple Gigabit Ethernet pipes at wire-rate. This is fueling the demand for 10 Gigabit Ethernet for NAS filer aggregation, as shown in Figure 5. In addition, there is growing demand for direct 10 Gigabit Ethernet connections to NAS filers to support high-performance applications that generate single data streams larger than 1 Gbps, which cannot be supported by 802.3ad link aggregation.

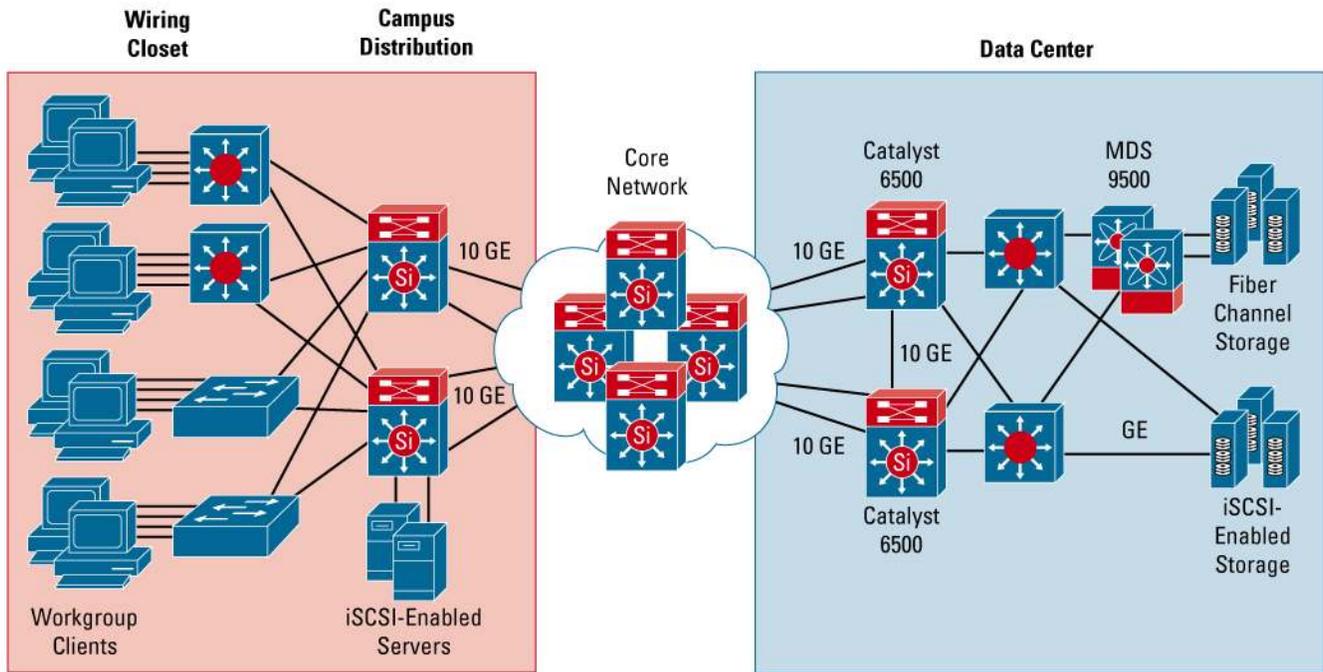
Besides providing high-performance access to shared files, a 10 Gigabit Ethernet infrastructure enables the added capability of filer-to-filer replication and backup to tape using protocols such as the Network Data Management Protocol (NDMP).

**Figure 5.** 10 Gigabit Ethernet for NAS Data Sharing and Consolidation



- Increasing Fan-Out to Shared Storage**—The rising costs of managing direct-attached storage, together with the growing capacity of storage subsystems to support hundreds of terabytes, is fueling the need to consolidate systems that were previously not considered as part of the storage network. The challenges in achieving this effectively center around the cost and scalability associated with extending Storage Area Networks (SANs) beyond a limited number of high-performance nodes. Enabling enterprise-wide access to storage over an IP network using the cost-effective iSCSI protocol is proving to be a very attractive way of achieving fan-out to the hundreds and thousands of servers that would otherwise be isolated from the storage network. In Figure 6, iSCSI-enabled servers in the campus can access the datacenter Fiber Channel SAN through the 10 Gigabit Ethernet infrastructure and the Cisco MDS 9500, which can act as an iSCSI gateway to Fiber Channel storage. 10 Gigabit Ethernet provides the network scalability needed to support the increasing number of distributed devices accessing shared storage across the enterprise.

**Figure 6.** 10 Gigabit Ethernet for Increasing Storage Consolidation Fan-Out

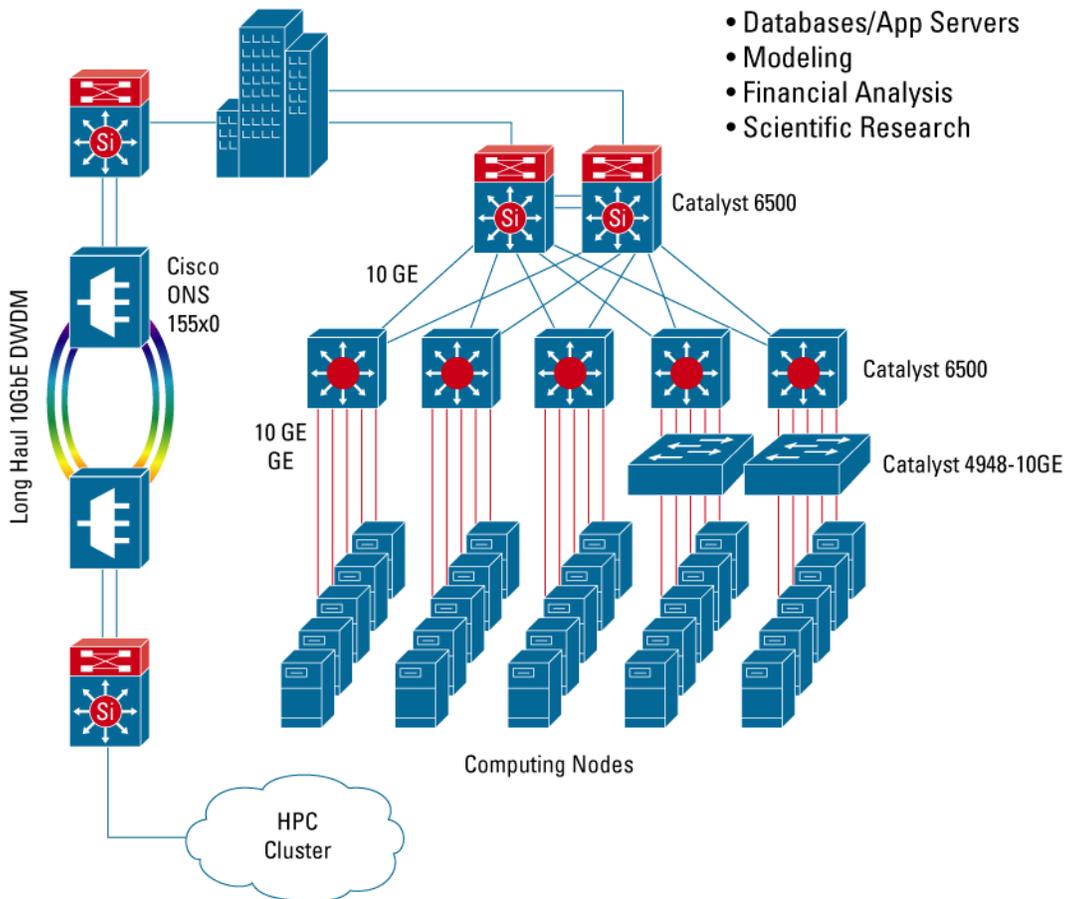


### Cluster and GRID Computing

Cluster and GRID computing is designed to meet the demands of CPU-intensive, transaction-intensive, and I/O-intensive applications that need more than a single server to efficiently complete the workload. Clustering provides a cost-effective way to scale computing needs beyond the confines of a single server and allows multiple computing nodes to work together as a large, virtual computing node. Cluster applications can be highly sensitive to the interconnect performance between computing nodes and thus place many demands on the networking infrastructure that link them together. Thus, clustered applications can benefit from the low-latency characteristics of 10 Gigabit Ethernet to maximize network performance. To significantly minimize server latency and CPU overhead, new server-side technologies are being introduced, such as system-level I/O Acceleration, TCP/IP Offload Engines (TOE), and Remote Direct Memory Access (RDMA). These major advancements in network and server performance also take advantage of the interoperability, management, and investment protection benefits of widely deployed Ethernet and IP technologies.

While clustered computing deployments have typically been used by the scientific research community, the commercial sector is increasingly using this paradigm. Database and application server vendors have added support for cluster computing in their products. Cluster computing is also being used for other high-performance computing (HPC) applications such as financial analysis and modeling, oil and gas exploration analysis, and engineering modeling. Figure 7 shows 10 Gigabit Ethernet providing high-performance connectivity within an HPC cluster and between distributed HPC clusters over DWDM.

**Figure 7.** 10 Gigabit Ethernet for Cluster and GRID Computing



**CONCLUSION**

10 Gigabit Ethernet deployments are rapidly growing as price and performance targets are met, new optics enable broader deployments, and the aggregate growth of new applications continue to increase bandwidth requirements. But 10 Gigabit Ethernet is just a network interface of a broader switching solution. Successful 10 Gigabit Ethernet deployments also incorporate leading intelligent switching services such as integrated security, high availability, delivery optimization, and enhanced manageability to provide the necessary support for new applications. In addition, to minimize costs, the transition to 10 Gigabit Ethernet should take advantage of existing switching investments in modules, chassis, and other components.

More information on the latest 10 Gigabit Ethernet technologies and Intelligent Switching services can be found at <http://www.cisco.com/go/switching>.

**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**

Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus  
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel  
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal  
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan  
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R) 205254.Y\_ETMG\_JR\_5.05

