



White Paper
Intel® PRO/10GbE
Server Adapters
Networking

10 Gigabit Ethernet Expands Network Bandwidth and Shrinks Latency

The newest evolutionary stage of Ethernet networking capability—10 Gigabit Ethernet or 10GbE—offers substantial additional capacity for handling the increasingly data-intensive, low-latency applications emerging in today's information technology environment. Such applications include use of 10GbE as a cost-effective fabric in network storage and for high-speed connectivity in grid computing and server farms. Additionally, 10GbE is appearing as a backbone in systems serving digital imaging, financial and scientific modeling, and other applications requiring fast transfer or sharing of massive amounts of data.

February 2006

Introduction

Contents

Introduction	2
The Need for 10 Gigabit Ethernet	3
Overview of 10GbE Technology	3
10 Gigabit Ethernet Applications	5
10GbE in Network Fabrics, Computing Clusters, and Related Applications	5
10GbE in Local Area Networks	7
10GbE in Metropolitan Area Networks	9
10GbE in Wide Area Networks	9
Performance-Enhancing Technologies for 10GbE Server Adapters	11
Conclusion	11

Since its inception in the mid 1970s, Ethernet has continuously evolved to increasingly higher performance levels. The original 10 Megabits-per-second (Mbps) bandwidth standard evolved to Fast Ethernet (100 Mbps), then to Gigabit Ethernet (GbE) and, most recently, to 10 Gigabit Ethernet (10GbE). And it is still evolving. The standards for 10GbE in fiber-optic media were ratified in 2002; the standard for 10GbE in CX4 copper transmission media was ratified in 2004; and ratification is expected in late 2006 for 10GbE in Category 5 (or better) unshielded twisted pair (UTP) copper.

Throughout its evolutionary history, Ethernet has remained unchanged in protocol compatibility across all versions. As a result, mixed-speed systems encompassing 10/100/1000 Mbps platforms and switching are commonplace. With regard to this history of compatibility, 10 Gigabit Ethernet is no different. It merges with existing Ethernet infrastructures to provide the higher-bandwidth, lower-latency capabilities needed to support the many data-intensive, fast-response applications emerging in today's financial, scientific, engineering and manufacturing communities.

This white paper provides an overview of the new 10 Gigabit Ethernet technology provided by Intel® PRO/10GbE Server Adapters and the major applications areas for 10GbE. Such applications include use of 10GbE as a cost-effective network fabric, as well as the uses and benefits of 10GbE in extending the range and capabilities of Local Area Networks (LANs), Metropolitan Area Networks (MANs) and Wide Area Networks (WANs).

The Need for 10 Gigabit Ethernet

Increasing network traffic (both in terms of the number of and size of transactions), migration of Gigabit Ethernet (GbE) to the desktop, and the growth of increasingly data-intensive applications are key factors behind the evolution to 10GbE. The higher bandwidth of 10GbE allows faster data flow to and from network servers, resulting in faster transaction servicing and completion, which frees the servers for handling more transactions faster.

The greater speed and faster response of 10GbE can benefit virtually any traffic-burdened network. However, 10GbE particularly benefits data-intensive applications. A good example is the creation of graphics-intensive animation films where increasingly detailed texturing, lighting and movement add tremendously to the file sizes with which animators collaboratively work. Without fast data exchange between workstations and servers, the level of realistic detail seen in today's films would be impossible. The same principle applies to a wide range of simulation and modeling applications, including enterprise financial applications, database and data-modeling simulations, weather forecasting, computer-aided design and manufacturing (CAD/CAM), and similar applications. Such applications demand fine-grained details and faster collaborative interactions, requiring higher bandwidth for faster transfer of larger files.

If bandwidth were the only issue in today's IT environment, questioning the need for 10GbE would be reasonable when link aggregation with multi-port GbE server adapters can provide nearly the same bandwidth benefit. Indeed, link aggregation with multi-port GbE adapters, such as the Intel® PRO/1000 Quad Port Server Adapter, may provide the best cost/benefit solution for performance enhancement in general-purpose networks, especially those that still rely exclusively on Category 5 copper cabling. However, 10GbE has the advantage of providing higher bandwidth in a single server adapter for single cable runs, which currently can be fiber-optic media or CX4 copper and, ultimately, will include Category 5 or better copper. Additionally, 10GbE offers low latency and longer cable runs in fiber-optic media—up to 10 kilometers (km) in single mode fiber, for example, with the Intel® PRO/10GbE LR Server Adapter.

Higher bandwidth and low latency in a single 10GbE server adapter are advantageous in a variety of data-center applications, such as grid computing, high-performance computing (HPC) clusters and wiring cabinet interconnects. Benefits also arise in using 10GbE to replace proprietary fabrics, such as Fibre Channel, that are used in Network Attached Storage (NAS). Not only can 10GbE provide the necessary connection performance, but also, because it is Ethernet rather than proprietary, 10GbE eliminates the specialized training and expense required for installing and maintaining proprietary fabrics.

The longer range of 10GbE in fiber-optic media offers some distinct advantages in the geographical organization and expansion of LANs. Campuses can be more widespread without resorting to special range-extending techniques, and data centers can be located more advantageously. For example, data centers can be placed between campuses in multi-campus organizations or they can be located off campus in lower-cost facilities. In essence, the LAN could actually be extended to become a MAN. Additionally, 10GbE connectivity can be extended from the LAN or MAN to Internet Points of Presence (PoPs) to implement a WAN capable of global connectivity for transmitting data at terabyte-per-hour rates.¹

All of these 10GbE applications, from data-center connectivity to global WANs, are discussed further in this white paper, along with the 10GbE technology that enhances their performance.

Overview of 10GbE Technology

The 10 Gigabit Ethernet Standard is an extension of the basic IEEE 802.3* standard protocols to a wire speed of 10 Gbps. As an extension, 10GbE is still fully Ethernet compatible and retains the key Ethernet architecture, including the Media Access Control (MAC) protocol, the Ethernet frame format, and the minimum and maximum frame size.

Just as Gigabit Ethernet followed the standard Ethernet model, 10GbE continues the evolution in speed while using virtually the same architecture used in other Ethernet specifications. This allows implementation of 10GbE while maintaining compatibility with the rest of the network architecture and retaining the existing principles of network operation and management.

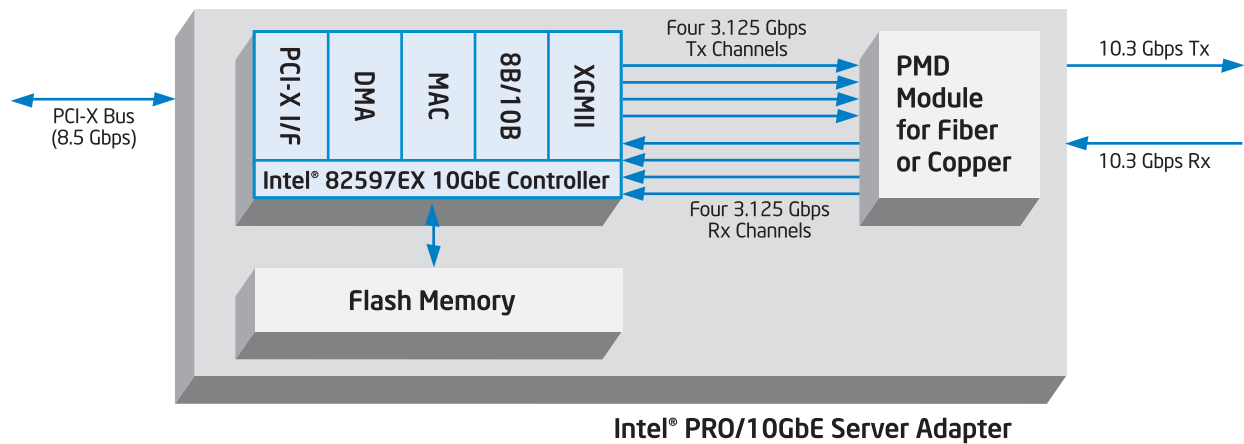


Figure 1. The architecture of the Intel® PRO/10GbE Server Adapters complies with the Ethernet standards and incorporates different Physical Medium Dependent modules for different transmission media.

However, some key differences exist in 10GbE. These differences, which are performance related and do not affect Ethernet compatibility, are illustrated in Figure 1. They involve primarily the MAC interface to the Physical Layer (PHY) of the Open Systems Interconnection (OSI) model and the Physical Medium Dependent (PMD) connection to the network transmission medium.

Also shown in Figure 1, the 10GbE Controller connects to the adapter's PMD by means of four transmit (Tx) channels and four receive (Rx) channels. Each channel has a 3.125 Gbps bandwidth so that the aggregate bandwidths in the Tx and Rx directions exceed 10Gbps each. Moreover, the Tx and Rx channels allow the server adapter and 10GbE traffic to operate in full-duplex mode (simultaneous Tx and Rx). By

contrast, Ethernet standards prior to 10GbE operated in simplex mode and required use of a carrier-sensing multiple-access with collision detection (CSMA/CD) protocol to resolve contention between network devices seeking access to the server. By using full-duplex operation and eliminating CSMA/CD processing, 10GbE can provide faster, lower-latency responses to network transaction requests.

The 10GbE standard also specifies various PMDs for connection to various transmission media. In reality, several versions of the 10GbE standard exist—with each version specifying a different medium, media connection (PMD type) and 10GbE Tx/Rx range for the medium. Table 1 lists the key versions of the 10GbE standard and their distinguishing attributes.

IEEE* Standard	Media	Distance	Server Adapter
802.3ae* 10GBase-LR	Single-Mode Fiber (1310 nm)	up to 10 km	Intel® PRO/10GbE LR
802.3ae 10GBase-SR	Multi-Mode Fiber (850 nm)	up to 300 m	Intel® PRO/10GbE SR
802.3ak* 10GBase-CX4	CX4 Copper Cable	up to 15 m	Intel® PRO/10GbE CX4
802.3an* 10GBase-T	Category 5 Copper UTP	up to 100 m (proposed)	Standard ratification expected in late 2006

Table 1. 10GbE Standards and Supporting Intel® PRO Server Adapters

10 Gigabit Ethernet Applications

Depending on existing network infrastructure or upgrade plans, the different Intel PRO/10GbE Server Adapters listed in Table 1 lend themselves to different application areas. For example, the Intel PRO/10GbE LR Server Adapter and Intel® PRO/10GbE SR Server Adapter fit well into existing fiber-based network structures. When planning to upgrade existing Category 5-based data centers and distribution links, single-mode fiber and the longer 10-km reach of the Intel PRO/10GbE LR Server Adapter is worth considering, especially for connecting widespread campus facilities or multiple campuses within a metropolitan area.

The Intel® PRO/10GbE CX4 Server Adapter is different, in that it provides CX4 copper cable connectivity. Because of the short range—15 meters—it is designed for short interconnects within the data center. It uses 8-pair, twin-axial InfiniBand*-style cabling and connectors, providing a cost-effective solution for high-bandwidth connectivity in server farms, network attached storage, grid computing, and other close-proximity data-center connections.

The following discussion expands further on 10GbE applications in the data center and in LANs, MANs and WANs. Note that, in addition to Ethernet compatibility, the Intel PRO/10GbE Server Adapters have been tested for interoperability with other networking components, such as 10GbE switches from Cisco and other major network component suppliers.

10GbE in Network Fabrics, Computing Clusters and Related Applications

Within the data center, 10GbE CX4 provides a low-cost, high-performance solution for numerous short-haul connectivity applications. These interconnect applications include:

- High-performance computing (HPC) clusters
- Consolidated LAN and iSCSI traffic
- Database links to backup storage and NAS
- Scaled-up, consolidated servers
- Real-time video streaming on Web servers

Outside of the data center, 10GbE CX4 can be used as a high-performance interconnect for workgroups involved in data-intensive, collaborative activities such as digital imaging and editing, simulation and modeling, and CAD/CAM. A typical case, as shown in Figure 2, would be the use of 10GbE CX4 for high-performance connectivity from a local workgroup server to the group's workstations, and 10GbE fiber connectivity for the longer-haul connection from the workgroup server to the production or manufacturing data center.

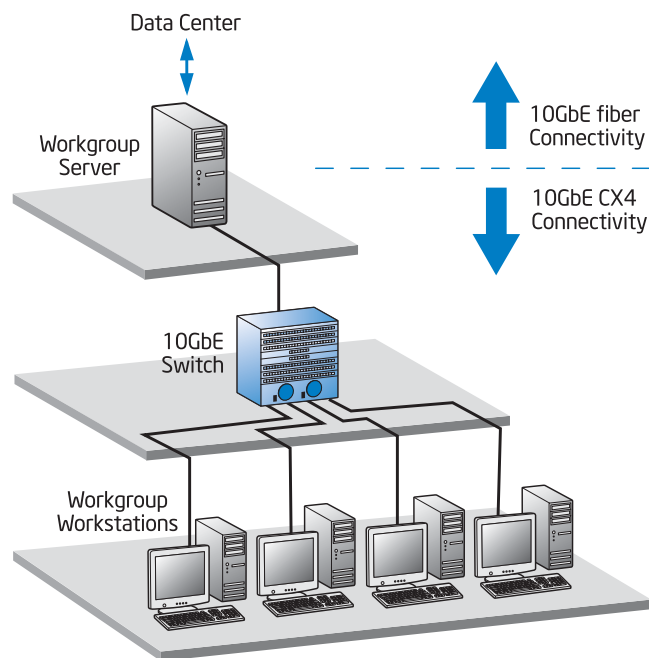


Figure 2. Example of 10GbE CX4 connectivity within a workgroup requiring high bandwidth and low latency for CAD/CAM, simulation and modeling, or other high-performance applications.

For these types of applications, the key benefit of 10GbE CX4 is high performance at a lower cost. On a per-port basis, 10GbE CX4 costs a fraction of 10GbE LR or SR fiber connectivity. On a per-Gigabit basis, 10GbE CX4 can also provide savings by consolidating teamed GbE adapters into a single 10GbE port. For example, to approach 10GbE performance, eight GbE ports need to be teamed. This can be done with minimum server slot consumption using two quad-port adapters for about the same adapter cost as a 10GbE CX4 adapter. However, the single cable and switch port for 10GbE CX4 operation saves over the additional cabling and switch ports used in eight-port adapter teaming. In addition, a power savings is realized in using one 10GbE CX4 adapter instead of teaming multiple GbE adapters. On the other hand, adapter teaming does offer the advantage of fault-tolerance through link failover capabilities.

Savings are also achieved when 10GbE CX4 is used instead of proprietary fabric interconnects in server area or storage area networks. Traditionally, proprietary fabrics such as InfiniBand, Servernet*, Myranet*, Wulkit* and Quadrics* have been the only way to meet the high-bandwidth, low-latency requirements of server and

storage area networks (SANs). However, such proprietary networks have a higher cost of ownership. Their initial cost is typically higher because the components tend to be low-volume, proprietary items. Installation and maintenance costs run higher, too, because proprietary fabrics require specialized skill sets. Also, interoperability problems often arise when mixing proprietary fabrics with standards-based network components. Such cost and interoperability issues can be avoided with 10GbE CX4. Because 10GbE CX4 is a non-proprietary, standards-based Ethernet technology, it is interoperable with other Ethernet components and does not require a special skill set for installation and maintenance.

Figure 3 illustrates the potential uses of 10GbE CX4 in the data center. These uses include application and database server clusters, and SAN connectivity provided by Intel® PRO/10GbE CX4 Server Adapters. Front-end or Web-server connectivity to the core can be either through multi-mode or single-mode fiber using the Intel PRO/10GbE SR or Intel PRO/10GbE LR Server Adapters, depending on the range needed to reach campus buildings.

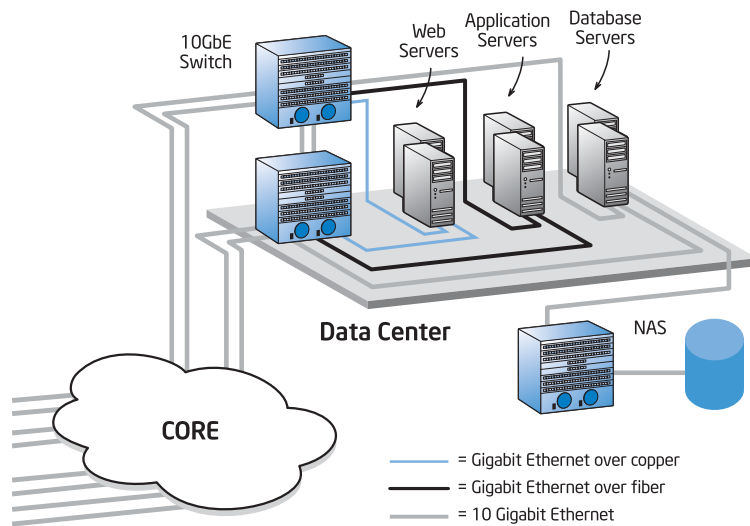


Figure 3. Examples of 10GbE connectivity in the data center. 10GbE CX4 is used for server cluster and network attached storage (NAS) connectivity, while the Intel® PRO/10GbE LR Server Adapter or Intel® PRO/10GbE SR Server Adapter provide fiber connectivity to the core.

10GbE in Local Area Networks

Use of 10GbE in a LAN configuration is shown in Figure 4. The data-center portion of this diagram is the same as discussed in Figure 3 and uses 10GbE CX4 connectivity for the database servers and storage. For campus buildings within 300 meters of the data center, multi-mode fiber and the Intel PRO/10GbE SR Server Adapter provide data-center connectivity through the core and to the two buildings.

For a more widespread campus, the 10 km range of the Intel PRO/10GbE LR Server Adapter and single-mode fiber are appropriate.

With 10GbE backbones, such as illustrated in Figure 4, organizations can easily support Gigabit Ethernet connectivity to workstations and desktops, resulting in much faster transaction servicing and completion and reduced network congestion. In addition to reducing congestion, 10GbE also reduces network latency because of its link speed, band-

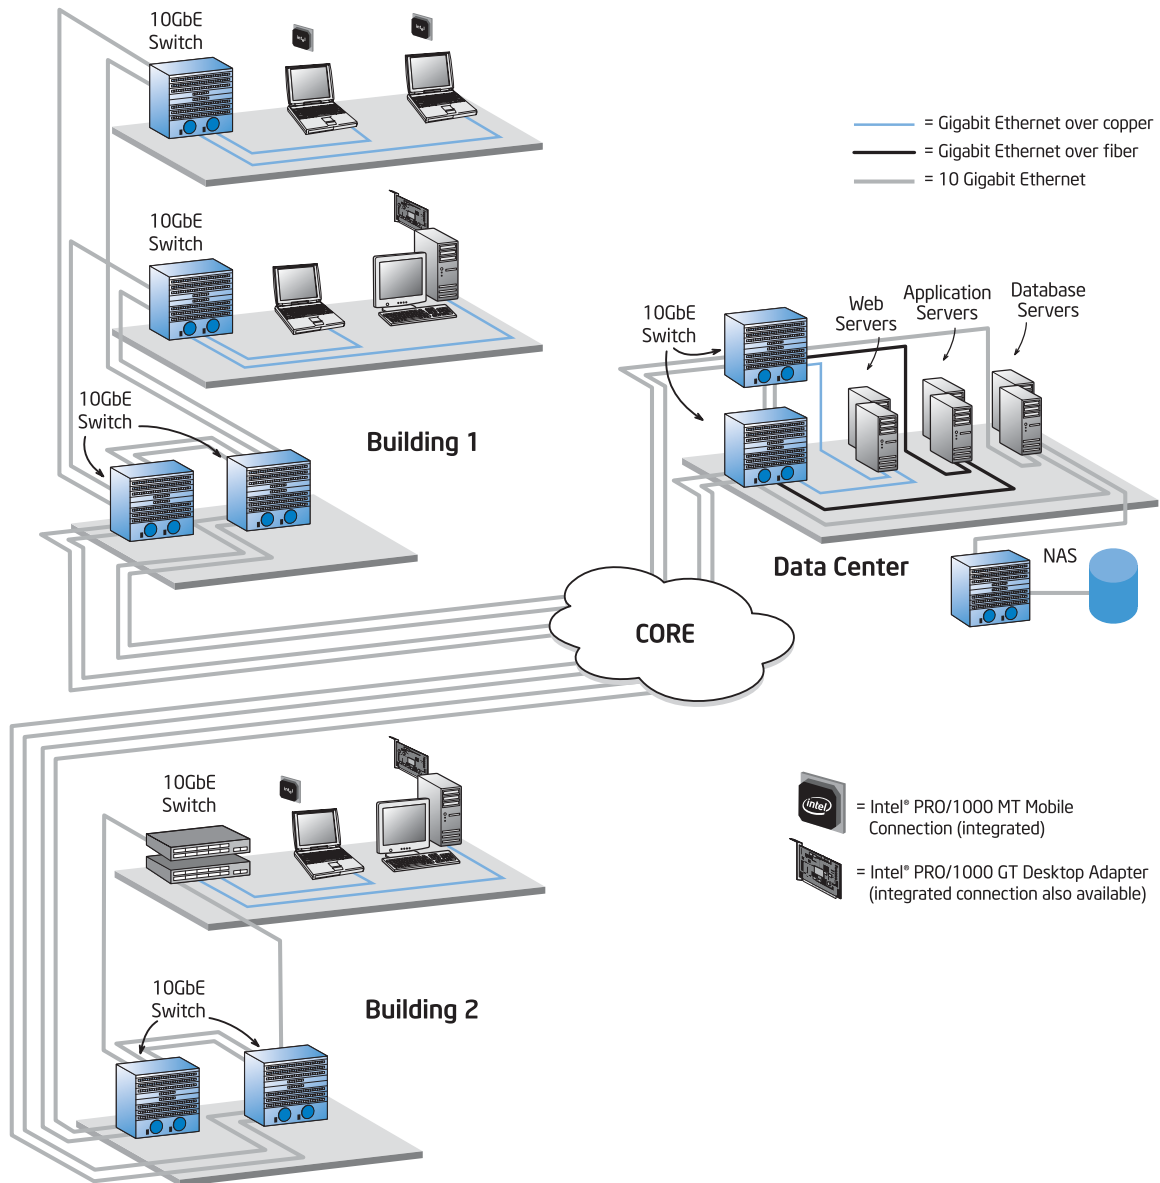


Figure 4. Example of using 10GbE in both the data center and LAN environments, with fiber media providing connectivity from the data center to the LAN.

width over-provisioning, and full-duplex operation that eliminates time-consuming CSMA/CD processing. The higher bandwidth and lower latency allow greater implementation of bandwidth-intensive applications throughout the network. Such applications include streaming video, interactive video conferencing, medical imaging, centralized applications, high-end graphics, and 3D simulation and modeling.

10GbE can also support geographic growth of the enterprise while providing greater flexibility in data center location. An example of this is shown in Figure 5, where Intel PRO/10GbE LR Server Adapters

and single-mode fiber are used to provide links of up to 10 km between a remotely located data center and local campus servers. The advantage here is that enterprise LAN growth is not confined to a single campus, but can spread out to more geographically or economically advantageous areas. Research and development, corporate headquarters, and production facilities can be placed at the most advantageous locations for each within the 10-km radius of reach provided by the Intel PRO/10GbE LR Server Adapter and single-mode fiber-optic media.

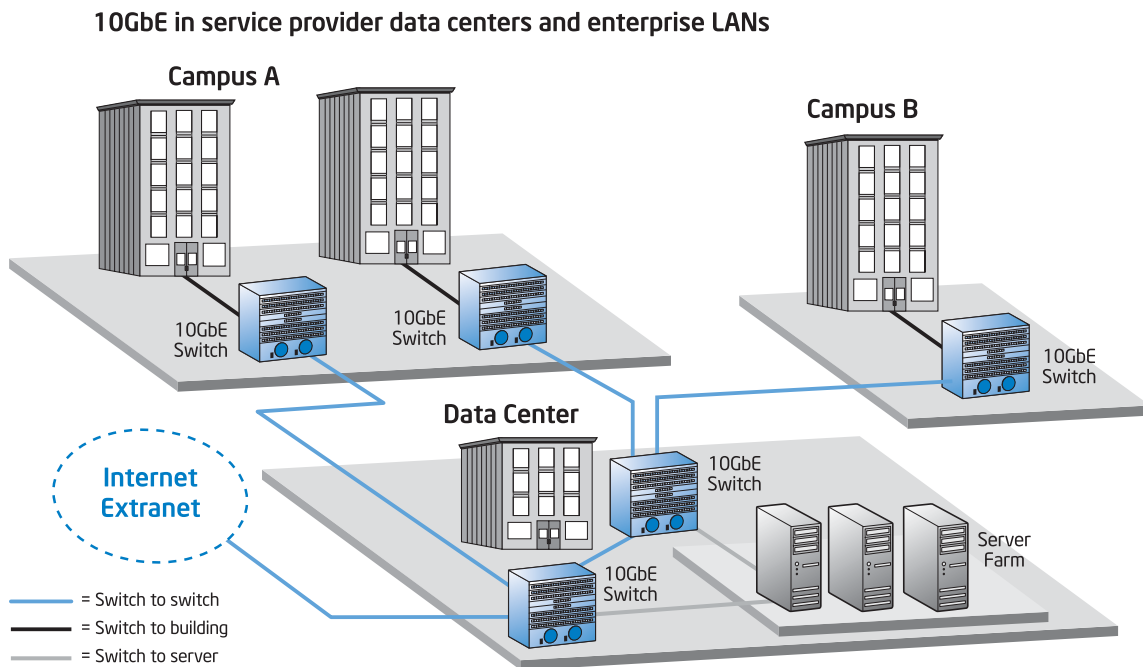


Figure 5. 10 Gigabit Ethernet in expanded LAN environments. The Intel® PRO/10GbE LR Server Adapter gives data centers a 10-km radius of reach in single-mode optical fiber.




Figure 6. Example of 10GbE in constructing MAN links. 10GbE can also play significant roles in storage infrastructure and server farm connectivity.

10GbE in Metropolitan Area Networks

Gigabit Ethernet is already in use as a backbone technology for dark-fiber MANs. With appropriate 10GbE interfaces, optical transceivers, and single-mode fiber, network designers and Internet service providers (ISPs) can now build links reaching 10 km or more, as indicated in Figure 6, to encircle metropolitan areas with city-wide networks.

Such MANs offer potential for more efficient and beneficial interaction between the public, private enterprises, educational institutions and governmental agencies. Enhanced telecommuting and video conferencing are often-cited benefits.

Interactive public education, access to land-use planning archives, and online archival research and data mining are just a few more possibilities that can benefit both the public and private sectors. To be truly effective, however, users must have rapid access to a large diversity of data in a variety of databases. In the past, the cost of storage devices and specialized NAS infrastructure proved to be cost prohibitive. That is no longer the case with today's significantly cheaper storage devices and the advent of 10 Gigabit Ethernet as a viable interconnect.

10 Gigabit Ethernet now enables cost-effective, high-speed infrastructure for both NAS and SAN. Prior to the introduction of

10GbE, some industry observers maintained that Ethernet lacked sufficient horsepower to get the job done. 10GbE now offers equivalent or superior data carrying capacity at latencies similar to many other storage networking technologies, including Fibre Channel, Ultra 160 or 320 SCSI, ATM OC-3, OC-12, and OC-192, and HIPPI (High-Performance Parallel Interface). Gigabit Ethernet storage servers, tape libraries and compute servers are already available; and 10GbE end-point devices are now available to complete the standards-based implementation of fast-access, network storage applications in MAN implementations.

10GbE in Wide Area Networks

The scientific community quickly perceived the potential benefits of 10GbE in WANs—and raced to prove those benefits as soon as 10GbE components became available. Shortly after the 2002 ratification of the IEEE 802.3ae 10GBase-LR standard, in February of 2003, a technical alliance consisting of Los Alamos National Laboratory (LANL), California Institute of Technology (CalTech), the Stanford Linear Accelerator Center (SLAC), and Conseil Européen pour la Recherche Nucléaire (CERN or European Organization for Nuclear Research) put 10GbE capabilities to test in a WAN of global proportions.

The technical alliance used the Intel PRO/10GbE LR Server Adapter to provide 10GbE links within the global WAN shown in Figure 7. This network spanned more than 9,600 km from Sunnyvale, California, to CERN in Geneva, Switzerland, and was able to carry more than a terabyte of data in less than an hour at a sustained rate of 2.38 Gbps.² This high data rate, while limited by the OC-48 transatlantic circuit, was still high enough to break the then-standing Internet2 Land Speed Record by 2.5 times. The accomplishment was noted in the *2004 Guinness Book of World Records*.

As always, records are meant to be broken, and the Internet2 Land Speed Record mark set in February 2003, using the Sunnyvale-to-Geneva WAN in Figure 7, has been broken again and again. Moreover, these record-breaking activities rely on one finely tuned connection pumping massive amounts of data, as opposed to a system handling multiple Transmission Control Protocol (TCP) connections. The point, however, is that the record-breaking performance invariably relies on 10GbE connectivity within the WAN. Even more important are the 10GbE proof-of-performance in WAN environments and the ability to provide high-bandwidth, low-latency transfers of massive data amounts across global boundaries. These provide important capabilities for global business and scientific pursuits. Examples

of such pursuits include CAD/CAM operations between facilities in different countries or continents, and global weather and oceanic data collection and modeling.

Poland's Poznań Supercomputing and Networking Center (PSNC) provides a practical example of 10GbE usage in a WAN environment with many TCP connections. PSNC operates a variety of networks, including the PIONIER network that has 6000 km of optical fiber connections with 21 Polish academic MANs and optical links to neighboring countries and the pan-European network, GÉANT.

Currently, PSNC's Clusterix project has linked approximately 800 Intel® Itanium® 2 processor-based systems across 12 distant organizations into a single grid-computing facility. Initially, multiple dedicated GbE connections provided the grid-computing links. However, PSNC found that multi-GbE links did not do well in such distributed grid-computing applications, particularly in handling bursty transmissions and because of degradation from packet reordering over diverse link paths. Feasibility studies using Intel PRO/10GbE LR Server Adapters revealed that the major performance issues, including packet reordering, could be resolved through dedicated 10GbE links. As a result, PSNC is migrating the Clusterix grid to 10GbE to achieve their high-performance, WAN-based grid-computing goals.

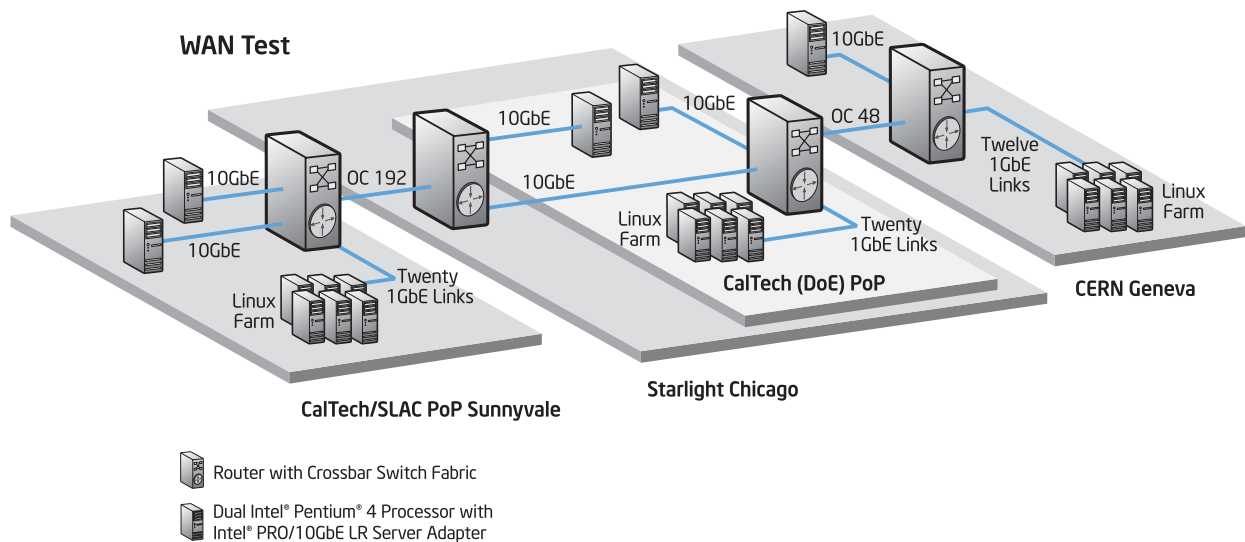


Figure 7. 10GbE in a WAN spanning 9,600 km from Sunnyvale, California to Geneva, Switzerland.

Performance-Enhancing Technologies for 10GbE Server Adapters

The 10GbE standards define the OSI Layer 1 and Layer 2 technologies and protocols necessary to achieve 10 Gigabit Ethernet performance and compatibility in various transmission media. Beyond that, a variety of technologies can be integrated into server adapters to enhance 10GbE performance. Typically, these performance-enhancing technologies increase throughput and reduce latency.

For example, the Intel PRO/10GbE Server Adapters currently incorporate the following performance-enhancing technologies:

- Jumbo frame support to 16 KB—allows use of larger data frames to reduce the amount of interrupts and interrupt processing required, thus increasing throughput and reducing latency.
- TCP/IP segmentation and TCP checksum offload in Rx and Tx—provides more efficient TCP/IP processing through the adapter and unburdens the host computer by performing checksum processing on the adapter.
- Mechanisms for delaying/reducing Tx and Rx frame interrupts—allows the adapter to perform interrupt moderation to reduce host CPU burden, resulting in more efficient overall traffic processing.

The above technologies—jumbo frame support, TCP/IP segmentation, TCP checksum offload, and reducing frame interrupts—provide on-adapter performance enhancements with the broadest applicability to today's high-performance network needs.

Conclusion

Ethernet is continuing its long history of evolution to meet increasing bandwidth demands in today's networks and applications. The newest evolutionary stage is 10 Gigabit Ethernet. 10GbE networking components are now playing a key role in providing high-bandwidth, low-latency capabilities in a wide variety of application areas.

The benefits of 10GbE include providing:


- A low-cost, Ethernet-compatible alternative to proprietary fabrics in high-performance computing clusters and network attached storage.
- Increased network performance and throughput in applications dealing with high traffic and large blocks of data, such as financial forecasting, CAD/CAM, data modeling and simulation, video editing, and computer generated imaging.
- Greater capacity and longer reach in LANs, MANs, and WANs.

To support the diversity of application needs and possibilities for 10GbE capabilities, Intel provides worldwide availability of the line of Intel® PRO/10GbE Server Adapters. These adapters and their key capabilities are as follows:

- **Intel PRO/10GbE CX4 Server Adapter**—provides CX4 copper cable connectivity up to 15 meters in server clusters, wiring cabinets, and network attached storage.
- **Intel PRO/10GbE SR Server Adapter**—provides multi-mode fiber connectivity up to 300 meters for data centers and LANs.
- **Intel PRO/10GbE LR Server Adapter**—provides single-mode fiber connectivity up to 10 kilometers for data centers, LANs, MANs, and WANs.

Availability of a 10GbE server adapter for Category 5 copper connectivity depends on completion and ratification of the 10GBase-T standard, expected in late 2006.

All Intel PRO/10GbE Server Adapters are designed and tested for full compliance with their respective 10GbE standards. Additionally, these adapters include a variety of performance enhancing technologies, such as jumbo frame support and TCP checksum offload, to increase throughput and reduce latency. All Intel PRO Server Adapters are also tested for interoperability with other network components from the world's major suppliers to ensure network connectivity that you can count on.



For more information on Intel® PRO/10GbE Server Adapters, visit
www.intel.com/go/10GbE

For information on the full line of Intel® PRO connectivity products, visit
www.intel.com/network

¹ http://www.intel.com/network/connectivity/case_studies/16832_LosAlamos_CS_r03.pdf

² "Optimizing 10 Gigabit Ethernet for Networks of Workstations, Clusters, and Grids: A Case Study" by Wu-chun Feng et al, *Proceedings of ACM/IEEE SC 2003: High-Performance Networking and Computing Conference*, November 2003, available at www.scconference.org/sc2003/paperpdfs/pap293.pdf or <http://public.lanl.gov/feng/sc03.pdf>.

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

Copyright ©2006 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel. Leap ahead. and Intel. Leap ahead. logo, Pentium and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.