Preparing Infrastructure for 10 Gigabit Ethernet Applications

Christopher T. Di Minico, Cable Design Technologies (CDT) Corporation,
Bruce Tolley, Cisco Systems

Background

The Internet will continue to grow, to change the way people communicate, to enable new modes of information in context and content, and to drive new higher bandwidth entertainment and business applications. To paraphrase the communication theorist Marshal McLuhan[1]: the Internet is the message; that is, the Internet has had a greater influence on the way we behave and interact than the information that is transmitted. McLuhan said, "We shape our tools, and they in turn shape us."
Two examples of this process can be observed in the impact to the music industry as a result of the Internet-enabled music file-swapping, and the erosion of literacy due to the short hand notations used in fast-paced Internet chat room sessions.

The Internet empowers people with the ability to instantly communicate data and voice content such as music, pictures, and videos, and to access information in new contexts such as e-mail, and online business and entertainment applications. The Internet has brought the electronic age to the consumer's doorstep as the milkman once brought the milk.

The escalating Internet data traffic has stimulated an unprecedented transformation in the telecommunication infrastructure. Network designers are struggling with critical path design choices to address a fundamental reengineering of the telecommunication infrastructure.

The TCP/IP protocol is the fundamental building block of the Internet. Since its beginning, TCP/IP has been associated with Ethernet, the most successful Layer 2 protocol in the history of networking. With over 300 million hub and switch ports installed, almost all Internet traffic today starts and ends on Ethernet attached workstations. Gigabit Ethernet requires laser sources for transmission over multimode and single mode fiber instead of light emitting diode sources (LED's). The laser-based Gigabit Ethernet technology now provides a viable alternative to other laser-based technologies;
operating distances and media types. Over 80 percent of Gigabit Ethernet ports today operate over optical fiber. Since February 2000, the IEEE 802.3ae Task Force has been developing specifications that increase the speed of Ethernet from 1 to 10 gigabits per second and support operation over optical fiber.

Introduction to 10 Gigabit Ethernet

With multiple media and technologies to choose from, the network upgrade path for the managers of enterprise and service provider networks are many and varied. Since March 1999, the Ethernet industry has been working on providing solutions to these problems by increasing the speed of Ethernet from 1 to 10 gigabits per second. For enterprise LAN applications, 10 Gigabit Ethernet will enable network managers to scale their Ethernet networks from 10 Mbps to 10,000 Mbps, while leveraging their investments in Ethernet as they increase their network performance. For service provider metropolitan and wide-area applications, 10 Gigabit Ethernet will provide high-performance, cost-effective links that are easily managed with Ethernet tools. 10 Gigabit Ethernet matches the speed of the fastest technology on the WAN backbone, OC-192, which runs at approximately 9.5 Gbps.

[1] Herbert Marshall McLuhan (1911-1980), Canadian writer, created a series of books about media, technology, and communications. McLuhan is best known for coining the phrase "the medium is the message," which became popular in the 1960s.

The IEEE 802.3ae 10 Gigabit Ethernet Task Force is chartered with developing the 10 Gigabit Ethernet specifications. This group is a subcommittee of the larger 802.3 Ethernet Working Group. The Task Force has achieved consensus not only on scope and purpose, but also on the core technical content. In contrast to previous Ethernet standards, 10 Gigabit Ethernet targets three application spaces: the LANs (including storage area networks), MANs, and WANs. While customers may see prestandard 10 Gigabit Ethernet products in 2001, the Task Force does not expect the standard to be completed until March 2002. A key milestone for the standards effort was the creation of Draft 3.0, which was forwarded to Working Group ballot in March 2001.

LAN Applications

10 Gigabit Ethernet has many potential applications for both service provider and enterprise networks. Figure 1 shows the standard LAN applications for 10 Gigabit Ethernet, which includes the following:

- Storage area networking (SAN) applications - Server interconnect for clusters of servers.
- Aggregation of multiple 1000BASE-X or 1000BASE-T segments into 10 Gigabit Ethernet downlinks.

- Switch-to-switch links for very high-speed connections between switches in the equipment room, in the same data center, or in different buildings.

Figure 1: 10 GbE LAN Applications

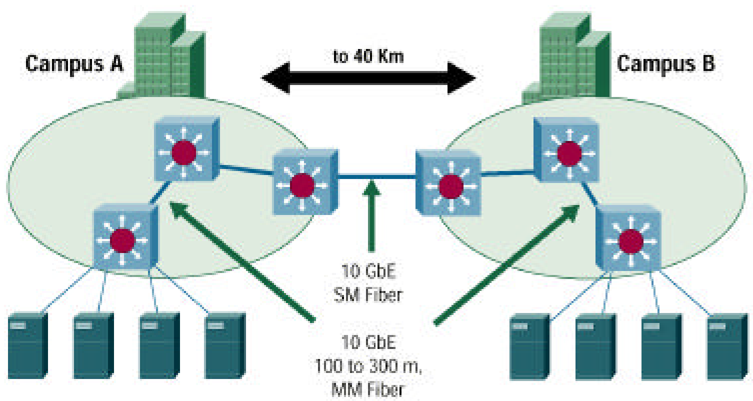

Source: Cisco Systems

Dark Fiber Metro Applications

One of the most exciting innovations in the Gigabit space has been the growth of the deployment of long distance Gigabit Ethernet using long wavelength optics on dark fiber to build network links that reach metropolitan distances.

10 Gigabit Ethernet, as a fundamental transport for facility services, will be deployed in MAN applications over dark fiber, and over dark wavelengths. The term "dark fiber" refers to unused single-mode fiber capacity from fiber that has been installed for long distance applications that usually reach up to 100 kilometers without amplifiers or optical repeaters. This fiber is not currently "lit," meaning that it is not carrying traffic and is not terminated to equipment.

As shown in Figure 2, 10 Gigabit Ethernet metropolitan networks will enable service providers to reduce the cost and complexity of their networks while increasing backbone capacity to 10 Gbps. This will be accomplished by eliminating the need to build out an infrastructure that contains not only several network elements required to run TCP/IP and data traffic, but also the network elements and protocols originally designed to transport voice. Reduction in the number of network elements and network layers lowers equipment costs, lowers operational costs, and simplifies the network architecture. With 10 Gigabit Ethernet backbone networks, service providers will be able to offer native 10/100/1000/10,000 Mbps Ethernet as a public service to customers, namely offering the customer twice the bandwidth of the fastest public MAN services OC-3 (155 Mbps) or OC-12 (622 Mbps) with no need for the added complexity of SONET or ATM, nor protocol conversions.

Figure 2: 10 GbE MAN Applications

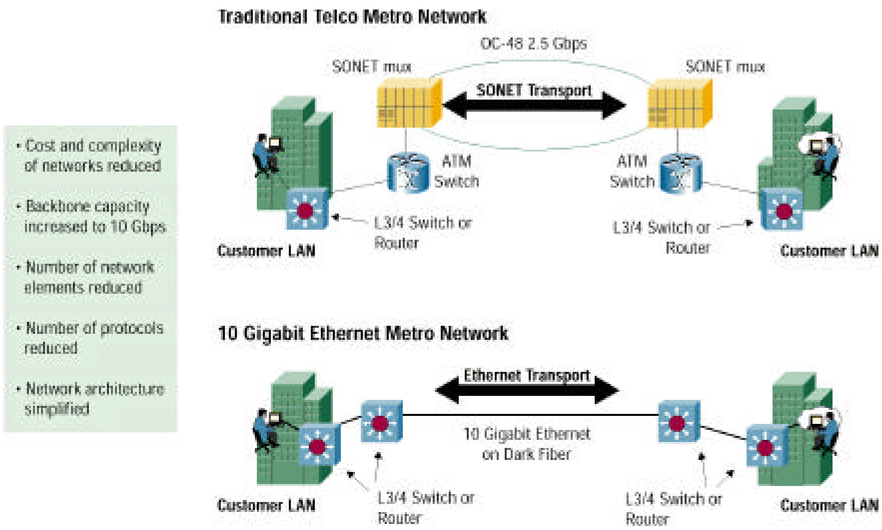

Source: Cisco Systems

Dark Wavelength Metro Applications with DWDM

As Figure 3 shows, 10 Gigabit Ethernet will be a natural fit to the dense wave division multiplexing (DWDM) equipment, which is deployed for metropolitan area applications. For enterprise networks, access to 10 Gigabit Ethernet services over DWDM will enable serverless buildings, remote backup, and disaster recovery. For service providers, 10 Gigabit Ethernet in the MAN will enable the provisioning of dark wavelength gigabit services at very competitive costs.

The terms “dark wavelength” or “dark lambda” refer to unused capacity available on a DWDM system. WDM is a long established technology in the WAN backbone that enables multiple data streams to be transformed into multiple, independent wavelengths. DWDM refers to systems that apply the tight wavelength spacing specified by the International Telecommunications Union (ITU), which is normally less than a nanometer (nm). Coarse or wide wavelength division multiplexing (CWDM or WWDM) refers to less costly optics that use wider spacing between wavelengths. The WDM device then multiplexes the multiple (16, 32, and 64) streams into one stream of “white light” across one fiber pair, increasing the bandwidth capacity of the link by a factor of 16, 32, or 64. At the opposite end, the multiple wavelengths are demultiplexed into the original data streams. Many MANs and much of the WAN backbone today contain installed DWDM equipment that has unused capacity or dark wavelengths.

Figure 3: 10 Gigabit Ethernet in the MAN over DWDM

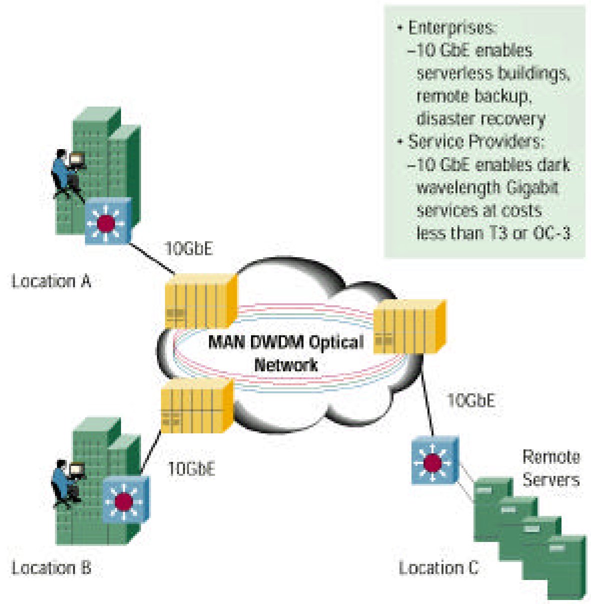

Source: Cisco Systems

10 Gigabit Ethernet WAN Applications

WAN applications for 10 Gigabit Ethernet look very similar to MAN applications: dark fiber, dark wavelength, and support for SONET infrastructure. Figure 4 shows multilayer switches and terabit routers attached via 10 Gigabit Ethernet to the SONET optical network, which includes add drop multiplexers (ADMs) and DWDM devices. When dark wavelengths are available, 10 Gigabit Ethernet can be transmitted directly across the optical infrastructure, reaching distances from 70 to 100 km. SONET is the dominant transport protocol in the WAN backbone today, and most MAN public services are offered as SONET OC-3 (155 Mbps) or OC-12 (622 Mbps).

Most of today's installed optical infrastructure is built out with a specific architecture and specific timing requirements to support OC-192 SONET. To make use of the SONET infrastructure, the IEEE 802.3ae Task Force specified a 10 Gigabit Ethernet interface (WAN PHY) that attaches to the SONET-based TDM access equipment at a data rate compatible with the payload rate of OC-192c/SDH VC-4-64c. This is accomplished by means of a physical layer link based on the WAN PHY between Gigabit or Terabit switches and Ethernet line-terminating equipment (LTE), which is attached to the SONET network. The WAN PHY interface does not attach directly to a SONET OC-192 interface.

The WAN PHY interface will allow the construction of MANs and WANs that connect geographically dispersed LANs between campuses or POPs through the SONET transport network. In other words, 10 Gigabit Ethernet interfaces that are compatible with SONET OC-192 payload rate facilitate the transport of native Ethernet packets across the WAN transport network, with no need for protocol conversion. Reducing the need for protocol conversion increases the performance of the network, makes it simpler and easier to manage, and less costly, because protocol conversion is CPU-intensive, adding complexity and additional elements to the network. Figure 4 shows that this WAN

PHY interface will enable Internet service providers (ISPs) and network service providers (NSPs) to create very high-speed links at very low cost between co-located, carrier-class switches and routers.

Figure 4: 10 Gigabit Ethernet in the WAN

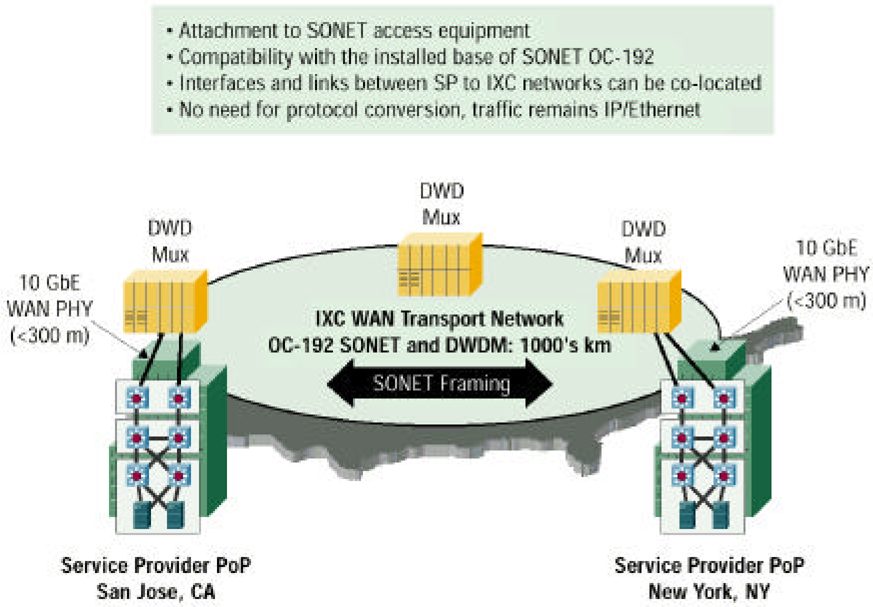

Source: Cisco Systems

10 Gbps Ethernet Physical Layer Specifications

The 10 Gigabit Ethernet physical layer specification, referred to as the “PHY”, provides the network manager and cabling distribution designer with the basic information required to select the appropriate optical transceiver types based on their network distance requirements, cabling performance, and types of network connections.

The 10 Gigabit Ethernet Standard defines two unique physical layer specifications associated with the types of network connections: the LAN physical layer (LAN PHY) and the WAN physical layer (WAN PHY). The physical layer (PHY) contains the types of transmitters and receivers and the functions that translate the data into signals (encoding), which are compatible with the cabling type used. The encoding function is performed in the physical coding layer (PCS) of the PHY. The LAN PHYs use 64B/66B encoded data; the WAN PHYs implement an encapsulation of the 64B/66B encoded data for compatibility with OC-192c/SDH VC-4-64c. The encapsulation is performed in the wide area network interface sublayer (WIS).

In Ethernet-speak, the transceiver types, which are cabling media dependent, are referred to as the physical media dependent (PMD) types. Examples of Ethernet optical fiber PMD types are 10BASE-F (10 Mbps), 100BASE-FX (100 Mbps), and 1000BASE-SX (1000 Mbps). The 10 Gigabit Ethernet PMDs include both serial and wavelength division multiplexing (WDM) fiber optic transceiver types. 10 Gigabit Ethernet operating distances are specified for both multimode and single mode fiber. The minimum operating distance for each option is associated with the targeted operating environment, i.e., LAN/MAN/WAN.

Serial PMD Implementations

The serial transmission is implemented by transmitting over one wavelength on each fiber. The nomenclature used to describe the serial PMD types include suffix designations which refer to both the wavelengths and the PHY types. The wavelength suffix designations are: S, L, and E, standing for short wavelength, long wavelength, and extra long wavelength. The designations R and W refer to the PHY types which are distinguished by the coding used. R stands for LAN PHY, which uses 64B/66B encoded data; W stands for WAN PHY, which implements an encapsulation of the 64B/66B encoded data for compatibility with OC-192c/SDH VC-4-64c. Table 1 and Table 2 provide descriptions for each serial PHY type.

Table 1. LAN PHY Type 10GBASE-R – PMD Types 10GBASE-SR/LR/ER

PMD	Description	PCS (10GBASE-R)
10GBASE-SR	850 nm short wavelength (10GBASE-S)	64B/66B encoding
10GBASE-LR	1310 nm long wavelength (10GBASE-L)	64B/66B encoding
10GBASE-ER	1550 nm extra long wavelength (10GBASE-E)	64B/66B encoding

Table 2. WAN PHY Type 10GBASE-W – PMD Types 10GBASE-SW/LW/EW

PMD	Description	WIS (10GBASE-W)
10GBASE-SW	850 nm short wavelength (10GBASE-S)	Encapsulated 64B/66B
10GBASE-LW	1310 nm long wavelength (10GBASE-L)	Encapsulated 64B/66B
10GBASE-EW	1550 nm extra long waveler (10GBASE-E)	Encapsulated 64B/66B

Tables 3-5 provide the wavelength, operating distance, and fiber type for each serial PMD. The serial options enable attachment to both multimode and single mode fibers. The LAN PHY and the WAN PHY will support the same cabling types and operating distances.

Table 3. 10GBASE-SR/SW

Fiber Type	Modal Bandwidth @850 nm (Mhz*km)	Minimum Range (meters)
62.5/125 µm MMF	160	2-28
62.5/125 µm MMF	200	2-35
50/125 µm MMF	400	2-69
50/125 µm MMF	500	2-86
50/125 µm MMF	2000	2-300
10 µm SMF	n/a	not supported

Table 4. 10GBASE-LR/LW

Fiber Type	Nominal wavelengt (nm)	Minimum Range (meters)
10 µm SMF	1310	2 – 10,000

Table 5. 10GBASE-ER/EW

Fiber Type	Nominal wavelengt (nm)	Minimum Range (meters)
10 µm SMF	1550	2 – 40,000

WWDM PMD Implementations

Advances in the control of laser sources and their packaging have enabled the launch of more than one wavelength into a single fiber. The transmission of multiple wavelengths on a single fiber is referred to as wavelength division multiplexing.

The nomenclature used to describe the WWDM PMD type 10GBASE-LX4 includes suffix designations which refer to both the wavelength and the PHY types. The wavelength suffix designation L stands for long wavelength. The suffix designation X stands for the encoding type, which is 8B/10B (8-bits of data into a 10-bit code groups). The suffix designation 4 refers to the number of wavelengths transmitted.

The 10GBASE-LX4 is implemented with a 1310 nm wide wavelength division multiplexing (WWDM) transceiver using two multimode or single mode fibers, transmitting and receiving over four separate wavelengths. Table 6 provides the wavelength, operating distance, and fiber type for the WWDM PMD type.

Table 6. 10GBASE-LX4

Fiber Type	Modal Bandw @1300 nm (minim OFL) (Mhz*km)	Minimum Range (Meters)
62.5/125 µm MMF	500	2-300
50/125 µm MMF	400	2-240
50/125 µm MMF	500	2-300
10 µm SMF	n/a	2-10,000

<u>10 Gigabit Ethernet network designs</u>

Key factors to consider in the design of 10 Gigabit Ethernet networks are:

- The fiber cabling type and the performance at a specified wavelength:
  - cabling attenuation
  - bandwidth

- The network topology, including operating distances and numbers of connectors, i.e., the optical link loss budget.

- The use of mode-conditioning patch cords if required. The 1310 nm WWDM solution, 10GBASE-LX4, requires the use of a mode-conditioning patch cord to achieve the range of operating distance of Table 6.

- The implementation of a cabling design, compatible with LED and laser-based Ethernet network devices, which will allow the integration of current LED based 10 Mbps and 100 Mbps networks and laser-based 1 Gbps and 10 Gbps networks.

<u>10 Gigabit Ethernet fiber optic link designs</u>

The 10 Gigabit Ethernet operating distances are limited by the channel insertion loss, the cable bandwidth for multimode, and the optical transceiver characteristics, i.e., PMD types.

Channel insertion loss is defined to account for cable and connector losses (Figure 5). The channel insertion loss is calculated using the specified cable loss for each operating distance and the loss of two connections. A connection consists of a mated pair of optical connectors. An allocation of 1.5 dB is budgeted for connections and splice losses.

Figure 5: Fiber Optic Cabling Channel

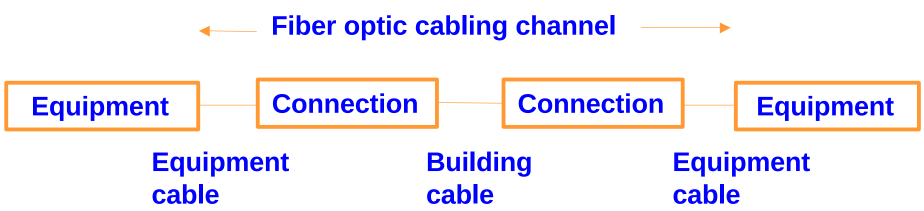

Source: CDT Corporation

The first step in the design of a fiber optic link is the characterization of the fiber optic power budget. The power budget is calculated by taking the difference between the minimum transmitter power launched into the fiber, and the receiver sensitivity. The receiver sensitivity is the minimum amount of power that is necessary to maintain the required signal-to-noise ratio over the specified operating conditions. The optical power budget determines the amount of total loss that can be introduced between the transmitter and the receiver.

For most networking applications, the power budget can be applied to characterize the loss available for fiber cable, fiber splices, and optical connectors. For 10 Gigabit Ethernet applications, the effects due to other signaling impairments (e.g., inter-symbol interference (ISI)) were accounted for by introducing additional power penalties. The difference between the power penalties and the channel insertion loss is the power budget margin. Figure 6 illustrates the optical link power budget.

Figure 6: Optical Link Power Budget

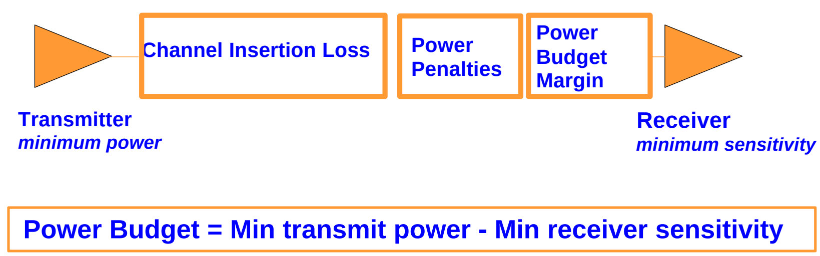

Source: CDT Corporation

Table 7provides the 10GBASE-LR/SW channel insertion loss, power budget, and power penalties.

Table 7. 10GBASE-LR/LW worst-case link power budget and penalties

PARAMETER	10GBASE-LR/LW	UNIT
Link power budget	9.4	dB
Operating distance	10	km
Channel insertion loss	7.17	dB
Allocation for penalties	2.96	dB

10 Gigabit Ethernet telecommunications cabling infrastructure

Ethernet has been traditionally associated with the LAN. Recognizing the need to support the predominately installed LAN cabling, the Ethernet standards have traditionally referenced the ANSI/TIA/EIA-568-A and ISO/IEC 11801 cabling standards. These standards specify telecommunications cabling within and between office oriented commercial buildings in a campus environment limited to a geographical extent of 3,000 m. Because these standards are limited in scope to traditional Ethernet LAN cabling topologies, they do not address the telecommunications cabling designs required for Ethernet in the WAN and MAN. In response to the need to support the evolution of Ethernet into the MAN and WAN, TIA-TR42, the engineering committee responsible for the development of ANSI/TIA/EIA-568, has initiated a study group to evaluate the cabling requirements for the 10 Gigabit Ethernet topologies.

10 Gigabit Ethernet applies to WAN and MAN topologies, including service provider backbones that are deployed between specialized networking distribution environments, such as Internet data centers, service provider distribution nodes, and central offices (Figure 7).The Internet data centers are designed to provide the network infrastructure and the physical environment necessary to house the active and passive networking components required for network accessibility 24 hours a day, 7 days a week. The service provider distribution nodes perform the traditional central office function of providing administration points where the service provider backbone cables are interconnected or cross-connected to networking equipment to facilitate circuit rearrangement or test access.

Figure 7. Ethernet Telecommunications Infrastructure

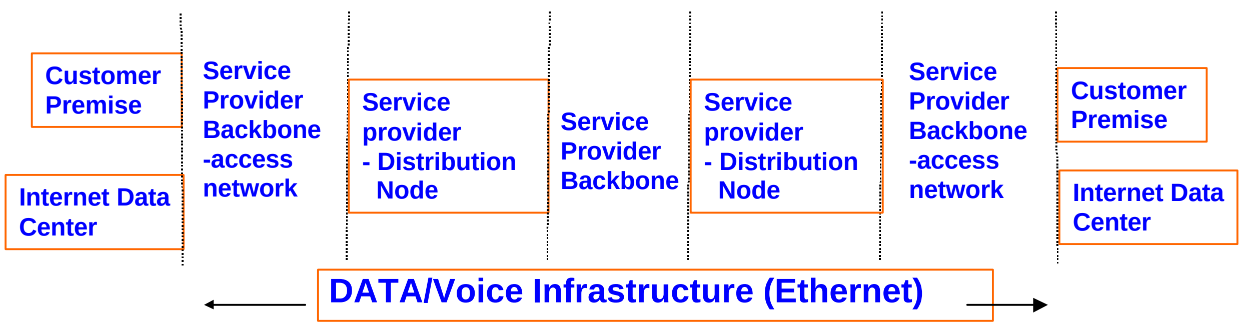

Source CDT Corporation

Conclusion

To conclude, 10 Gigabit Ethernet might well become the technology of choice for enterprise, metropolitan, and wide area networks. In terms of physical media, 10 Gigabit Ethernet will support distances to 300 meters on multimode fiber and 40 km or more on single mode fiber. With 10 Gigabit Ethernet, enterprise network managers and service providers will be able to build LANs, MANs, and WANs using Ethernet as the end-to-end Layer 2 transport. Long-distance reach on single mode fiber enables enterprise network managers and service providers to build simple, low-cost, metropolitan-sized networks with Layer 3-4 switches and 10 Gigabit Ethernet backbones. In addition, 10 Gigabit Ethernet will support an optional SONET-friendly PHY to enable transmission of Ethernet over the SONET transport infrastructure.

Authors

Mr. Christopher T. Di Minico is Vice President and Chief Technology Officer for the Cable Design Technologies Corporation (CDT) Network Divisions. Mr. Di Minico has over 30 years of experience in the telecommunication industry. Prior to his current position with CDT, Mr. Di Minico was employed at Digital Equipment Corporation with responsibility for the development and delivery of telecommunication connectivity products and services. Mr. Di Minico plays an active role in the development of Telecommunication Standards as a Member of the Institute of Electrical and Electronic Engineers (IEEE), the Telecommunications Industry Association (TIA), and the US advisory group for international cabling standards development. He is the elected liaison for the TIA TR42 committee to IEEE 802.3.

At Cisco Systems **Mr. Bruce Tolley** is Manager, Emerging Technologies with responsibility for 10 Gigabit Ethernet networking in the Gigabit Systems Business Unit. He is also Vice President and a Director of the 10 Gigabit Ethernet Alliance, of which Cisco is a founding member. Mr. Tolley also represents Cisco in the IEEE 802.3 Ethernet in the First Mile (EFM) Study Group, the IEEE 802.3ae 10 Gigabit Ethernet Task Force, and the 10 Gigabit Ethernet Alliance Technical Committee. Prior to his current position, Mr. Tolley held various business development and marketing management positions at 3Com Corporation.

Mr. Tolley is a frequent speaker on high-speed Ethernet networking at conferences such as Networld+Interop and Next Generation Networks and has published trade press articles including “Get Ready for 10 Gigabit Ethernet, Network World 6/26/00,” ”Running 1000BASE-T:Gigabit Ethernet Over Copper,” Gigabit Ethernet Alliance, May 1999, “Link Aggregation: Bolstering Bandwidth," Network World, 11/30/98, and "Stop Bashing ATM," Solutions Integrator Magazine, IDG Publications, November 1997. Mr. Tolley studied at Cambridge University, U.K. and Tuebingen University, Germany and holds post-graduate degrees from Stanford University and an MBA from the Haas School of Business, U.C. Berkeley.