CHAPTER 1

Introducing DWDM

The following discussion provides some background on why dense wavelength division multiplexing (DWDM) is an important innovation in optical networks and what benefits it can provide. We begin with a high-level view of the segments of the global network and the economic forces driving the revolution in fiber optic networks. We then examine the differences between traditional time-division multiplexing (TDM) and wavelength division multiplexing (WDM). Finally, we explore the advantages of this new technology.

Global Network Hierarchy

It is the nature of modern communications networks to be in a state of ongoing evolution. Factors such as new applications, changing patterns of usage, and redistribution of content make the definition of networks a work in progress. Nevertheless, we can broadly define the larger entities that make up the global network based on variables such as transport technology, distance, applications, and so on.

One way of describing the metropolitan area network (MAN) would be to say that it is neither the long-haul nor the access parts of the network, but the area that lies between those two (see Figure 1-1).

Figure 1-1 Global Network Hierarchy

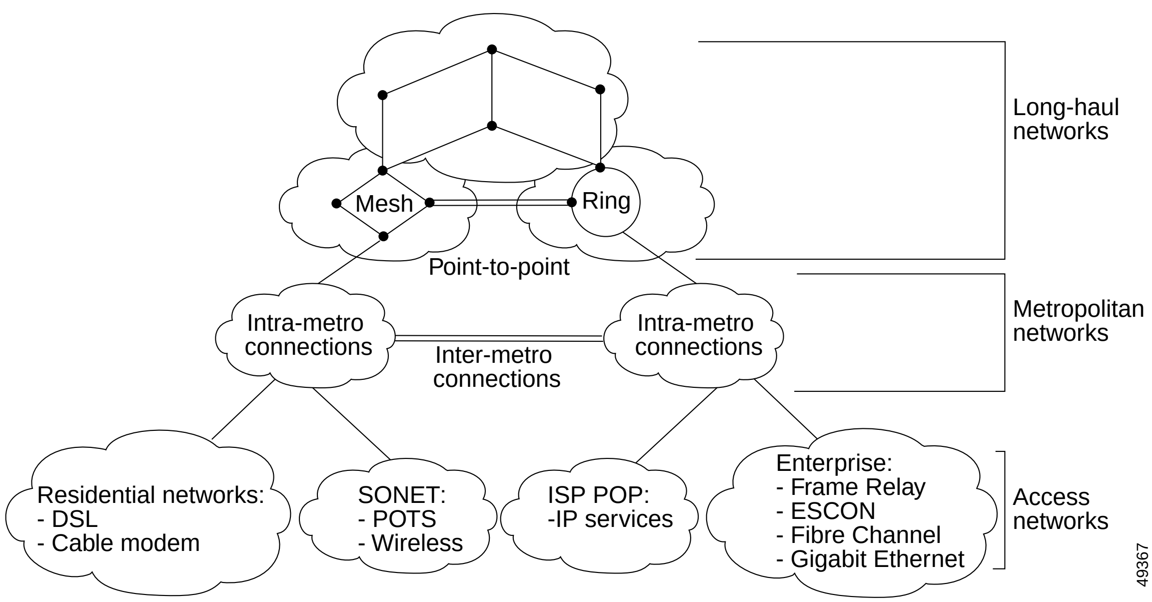

Long-Haul Networks

Long-haul networks are at the core of the global network. Dominated by a small group of large transnational and global carriers, long-haul networks connect the MANs. Their application is transport, so their primary concern is capacity. In many cases these networks, which have traditionally been based on Synchronous Optical Network (SONET) or Synchronous Digital Hierarchy (SDH) technology, are experiencing fiber exhaust as a result of high bandwidth demand.

Access Networks

At the other end of the spectrum are the access networks. These networks are the closest to the end users, at the edge of the MAN. They are characterized by diverse protocols and infrastructures, and they span a broad spectrum of rates. Customers range from residential Internet users to large corporations and institutions. The predominance of IP traffic, with its inherently bursty, asymmetric, and unpredictable nature, presents many challenges, especially with new real-time applications. At the same time, these networks are required to continue to support legacy traffic and protocols, such as IBM's Enterprise System Connection (ESCON).

Metropolitan Area Networks

Between these two large and different networking domains lie the MANs. These networks channel traffic within the metropolitan domain (among businesses, offices, and metropolitan areas) and between large long-haul points of presence (POPs). The MANs have many of the same characteristics as the access networks, such as diverse networking protocols and channel speeds. Like access networks, MANs have been traditionally SONET/SDH based, using point-to-point or ring topologies with add/drop multiplexers (ADMs).

The MAN lies at a critical juncture. On the one hand, it must meet the needs created by the dynamics of the ever-increasing bandwidth available in long-haul transport networks. On the other hand, it must address the growing connectivity requirements and access technologies that are resulting in demand for high-speed, customized data services.

Metropolitan and Long-Haul Networks Compared

There is a natural tendency to regard the MAN as simply a scaled-down version of the long-haul network. It is true that networks serving the metropolitan area encompass shorter distances than in the long-haul transport networks. Upon closer examination, however, these differences are superficial compared to other factors. Network shape is more stable in long-haul, while topologies change frequently in the MAN. Many more types of services and traffic types must be supported in MANs, from traditional voice and leased line services to new applications, including data storage, distributed applications, and video. The long-haul, by contrast, is about big pipes.

Another important way in which metropolitan networks today differ from trunk-oriented long haul networks is that they encompass a collection of low bit-rate asynchronous and synchronous transmission equipment, short loops, small cross-sections, and a variety of users with varying bandwidth demands. These fundamental differences between the two types of networks have powerful implications for the requirements in the metropolitan domain. Protocol and speed transparency, scalability, and dynamic provisioning are at least as important as capacity, which rules in the long-haul market.

An Alternative View

The preceding breakdown of the global network represents a somewhat simplified view. In reality, the lines between the domains are not always so clear-cut. Long-haul and metropolitan networks are sometimes not clearly delineated; the same holds true for the access and metropolitan domains.

Furthermore, other views of the global network exist. One, for example, defines the access network as part of, rather than separate from, the MAN, while also including enterprise connectivity in the metropolitan domain. In this view, the metropolitan market breaks down as follows:

- Core—These are essentially scaled-down long-haul systems. They are considered the core of the MAN, because they interconnect carrier POPs and do not directly interface with end users.
- Metropolitan access—This is the segment between carrier POPs and access facilities, which could be equipment at customer premises or at an aggregation point.
- Enterprise—This is the part of the network dedicated to serving the needs of enterprises. Using owned or leased fiber (or leased fiber capacity), connectivity is provided between geographically disparate enterprise sites and for new applications, such as storage area networks (SANs).

Economic Forces

As we enter the twenty-first century, it goes without saying that information services have permeated society and commerce. Information, while still a tool, has become a commodity in itself. Yet the universal acceptance and ubiquitous adoption of information technology systems has strained the backbones on which they were built. High demand—coupled with high usage rates, a deregulated telecommunications environment, and high availability requirements—is rapidly depleting the capacities of fibers that, when installed 10 years ago, were expected to suffice for the foreseeable future.

Bandwidth Demand

The explosion in demand for network bandwidth is largely due to the growth in data traffic, specifically Internet Protocol (IP). Leading service providers report bandwidths doubling on their backbones about every six to nine months. This is largely in response to the 300 percent growth per year in Internet traffic, while traditional voice traffic grows at a compound annual rate of only about 13 percent (see Figure 1-2).

Figure 1-2 Data Traffic Overtakes Voice Traffic

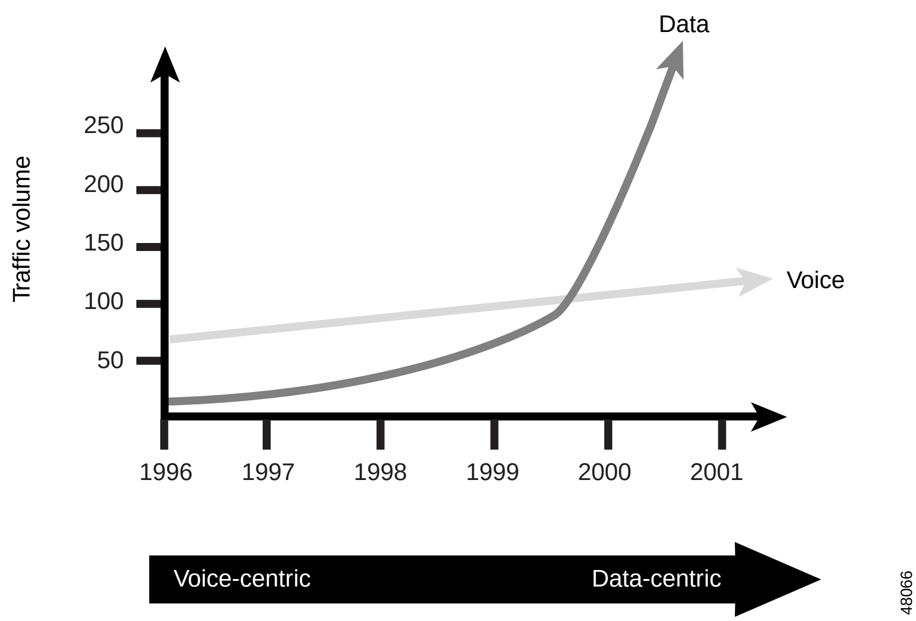

At the same time that network traffic volume is increasing, the nature of the traffic itself is becoming more complex. Traffic carried on a backbone can originate as circuit based (TDM voice and fax), packet based (IP), or cell based (ATM and Frame Relay). In addition, there is an increasing proportion of delay sensitive data, such as voice over IP and streaming video.

In response to this explosive growth in bandwidth demand, along with the emergence of IP as the common foundation for all services, long-haul service providers are moving away from TDM based systems, which were optimized for voice but now prove to be costly and inefficient. Meanwhile, metropolitan networks are also experiencing the impact of growing congestion, as well as rapidly changing requirements that call for simpler and faster provisioning than is possible with older equipment and technologies. Of key importance in the metropolitan area is the growth in storage area networks (SANs), discussed in the "Storage Area Networks" section on page 3-5.

Competition and Reliability

While the demand for bandwidth is driven largely by new data applications, Internet usage, and the growth in wireless communications, two additional factors come into play: competition and network availability.

The telecommunication sector, long a beneficiary of government regulation, is now a highly competitive industry. Competition was first introduced into the U.S. long-distance market in 1984, and the 1996 Telecommunications Reform Act is now resulting in an increasingly broad array of new operators. These new carriers are striving to meet the new demand for additional services and capacity.

There are two main effects on the industry from competition:

- Enhanced services are created by newcomers trying to compete with incumbents. In the metropolitan market, for example, there are broadband wireless and DSL services to homes and small and medium-sized business, high-speed private line and VPN services to corporations, and transparent LAN services to enterprise network customers.
- New carriers coming onto the scene create new infrastructure so that they do not have to lease from existing operators. Using this strategy, they have more control over provisioning and reliability.

As telecommunications and data services have become more critical to business operations, service providers have been required to ensure that their networks are fault tolerant. To meet these requirements, providers have had to build backup routes, often using simple 1:1 redundancy in ring or point-to-point configurations. Achieving the required level of reliability, however, means reserving dedicated capacity for failover. This can double the need for bandwidth on an already strained infrastructure (see Figure 1-3).

Figure 1-3 Reserving Bandwidth Reduces Overall Capacity

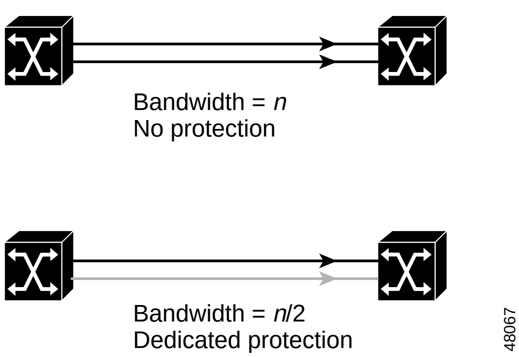

Options for Increasing Carrier Bandwidth

Faced with the challenge of dramatically increasing capacity while constraining costs, carriers have two options: Install new fiber or increase the effective bandwidth of existing fiber.

Laying new fiber is the traditional means used by carriers to expand their networks. Deploying new fiber, however, is a costly proposition. It is estimated at about $70,000 per mile, most of which is the cost of permits and construction rather than the fiber itself. Laying new fiber may make sense only when it is desirable to expand the embedded base.

Increasing the effective capacity of existing fiber can be accomplished in two ways:

- Increase the bit rate of existing systems.
- Increase the number of wavelengths on a fiber.

Increase the Bit Rate

Using TDM, data is now routinely transmitted at 2.5 Gbps (OC-48) and, increasingly, at 10 Gbps (OC-192); recent advances have resulted in speeds of 40 Gbps (OC-768). The electronic circuitry that makes this possible, however, is complex and costly, both to purchase and to maintain. In addition, there are significant technical issues that may restrict the applicability of this approach. Transmission at OC-192 over single-mode (SM) fiber, for example, is 16 times more affected by chromatic dispersion than the next lower aggregate speed, OC-48. The greater transmission power required by the higher bit rates also introduces nonlinear effects that can affect waveform quality. Finally, polarization mode dispersion, another effect that limits the distance a light pulse can travel without degradation, is also an issue. These characteristics of light in fiber are discussed further in the "Optical Fibers" section on page 2-5.

Increase the Number of Wavelengths

In this approach, many wavelengths are combined onto a single fiber. Using wavelength division multiplexing (WDM) technology several wavelengths, or light colors, can simultaneously multiplex signals of 2.5 to 40 Gbps each over a strand of fiber. Without having to lay new fiber, the effective capacity of existing fiber plant can routinely be increased by a factor of 16 or 32. Systems with 128 and 160 wavelengths are in operation today, with higher density on the horizon. The specific limits of this technology are not yet known.

Time-Division Multiplexing

Time-division multiplexing (TDM) was invented as a way of maximizing the amount of voice traffic that could be carried over a medium. In the telephone network before multiplexing was invented, each telephone call required its own physical link. This proved to be an expensive and unscalable solution. Using multiplexing, more than one telephone call could be put on a single link.

TDM can be explained by an analogy to highway traffic. To transport all the traffic from four tributaries to another city, you can send all the traffic on one lane, providing the feeding tributaries are fairly serviced and the traffic is synchronized. So, if each of the four feeds puts a car onto the trunk highway every four seconds, then the trunk highway would get a car at the rate of one each second. As long as the speed of all the cars is synchronized, there would be no collision. At the destination the cars can be taken off the highway and fed to the local tributaries by the same synchronous mechanism, in reverse.

This is the principle used in synchronous TDM when sending bits over a link. TDM increases the capacity of the transmission link by slicing time into smaller intervals so that the bits from multiple input sources can be carried on the link, effectively increasing the number of bits transmitted per second (see Figure 1-4).

Figure 1-4 TDM Concept

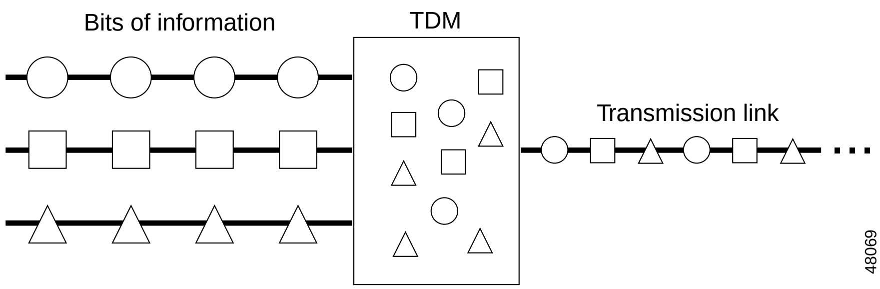

With TDM, input sources are serviced in round-robin fashion. Though fair, this method results in inefficiency, because each time slot is reserved even when there is no data to send. This problem is mitigated by the statistical multiplexing used in Asynchronous Transfer Mode (ATM). Although ATM offers better bandwidth utilization, there are practical limits to the speed that can be achieved due to the electronics required for segmentation and reassembly (SAR) of ATM cells that carry packet data.

SONET and TDM

The telecommunications industry adopted the Synchronous Optical Network (SONET) or Synchronous Digital Hierarchy (SDH) standard for optical transport of TDM data. SONET, used in North America, and SDH, used elsewhere, are two closely related standards that specify interface parameters, rates, framing formats, multiplexing methods, and management for synchronous TDM over fiber.

SONET/SDH takes *n* bit streams, multiplexes them, and optically modulates the signal, sending it out using a light emitting device over fiber with a bit rate equal to (incoming bit rate) x *n*. Thus traffic arriving at the SONET multiplexer from four places at 2.5 Gbps will go out as a single stream at 4 x 2.5 Gbps, or 10 Gbps. This principle is illustrated in Figure 1-5, which shows an increase in the bit rate by a factor of four in time slot *T*.

Figure 1-5 SONET TDM

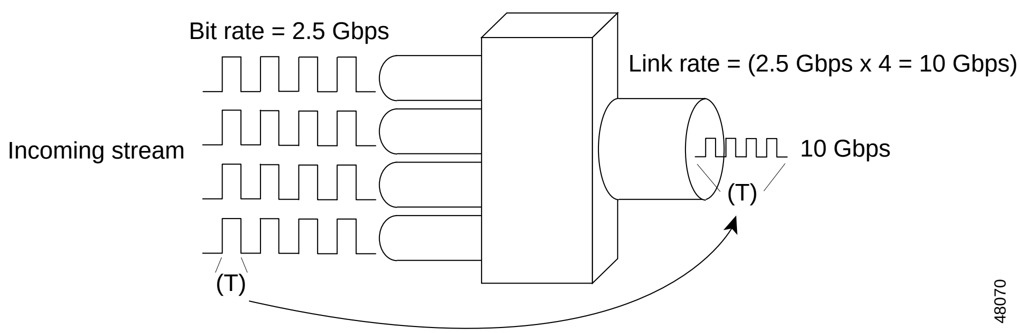

The original unit used in multiplexing telephone calls is 64 kbps, which represents one phone call. Twenty-four (in North America) or thirty-two (outside North America) of these units are multiplexed using TDM into a higher bit-rate signal with an aggregate speed of 1.544 Mbps or 2.048 Mbps for transmission over T1 or E1 lines, respectively.The hierarchy for multiplexing telephone calls is shown in Table 1-1.

Table 1-1 Telco Multiplexing Hierarchy

Signal	Bit Rate	Voice Slots
DS0	64 kbps	1 DS0
DS1	1.544 Mbps	24 DS0s
DS2	6.312 Mbps	96 DS0s
DS3	44.736 Mbps	28 DS1s

These are the basic building blocks used by SONET/SDH to multiplex into a standard hierarchy of speeds, from STS-1 at 51.85 Mbps to STS-192/STM-64 at 10 Gbps. Table 1-2 shows the relationship between the telco signal rates and the most commonly used levels of the SONET/SDH hierarchy (OC-768 is not yet common).

Table 1-2 SONET/SDH Multiplexing Hierarchy

Optical Carrier	SONET/SDH Signal	Bit Rate	Capacity
OC-1	STS-1	51.84 Mbps	28 DS1s or 1 DS3
OC-3	STS-3/STM-1	155.52 Mbps	84 DS1s or 3 DS3s
OC-12	STS-12/STM-4	622.08 Mbps	336 DS1s or 12 DS3s
OC-48	STS-48/STM-16	2488.32 Mbps	1344 DS1s or 48 DS3s
OC-192	STS-192/STM-64	9953.28 Mbps	5379 DS1s or 192 DS3s

Figure 1-6 depicts this multiplexing and aggregation hierarchy. Using a standard called virtual tributaries for mapping lower-speed channels into the STS-1 payload, the 28 DS1 signals can be mapped into the STS-1 payload, or they can be multiplexed to DS3 with an M13 multiplexer and fit directly into the STS-1. Note also that ATM and Layer 3 traffic, using packet over SONET (POS), can feed into the SONET terminal from switches equipped with SONET interfaces.

Figure 1-6 TDM and SONET Aggregation

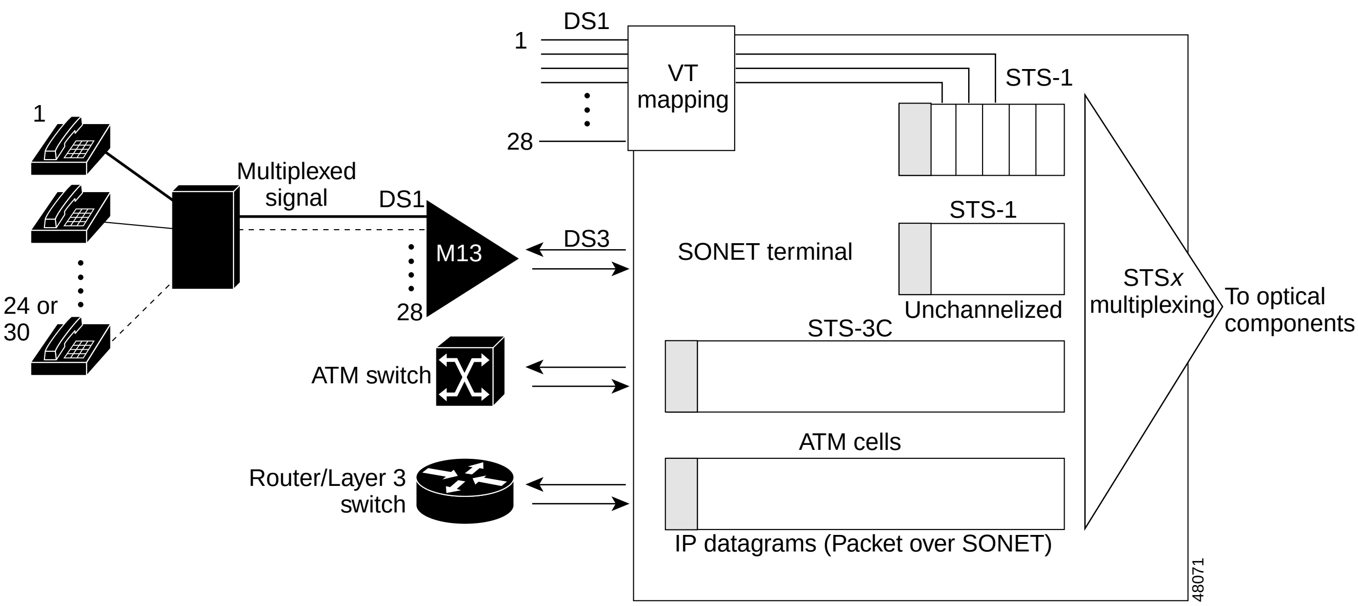

SONET/SDH does have some drawbacks. As with any TDM, the notions of priority or congestion do not exist in SONET or SDH. Also, the multiplexing hierarchy is a rigid one. When more capacity is needed, a leap to the next multiple must be made, likely resulting in an outlay for more capacity than is initially needed. For example, the next incremental step from 10 Gbps (STS-192) TDM is 40 Gbps (STS-768). Also, since the hierarchy is optimized for voice traffic, there are inherent inefficiencies when carrying data traffic with SONET frames. Some of these inefficiencies are shown in Table 1-3. DWDM, by contrast, can transport any protocol, including SONET, without special encapsulation.

Table 1-3 Ethernet in SONET Inefficiencies

Ethernet	SONET/SDH Signal	Bit Rate	Wasted Bandwidth
10BASE-T (10 Mbps)	STS-1	51.8540 Mbps	80.709%
100BASE-T (100 Mbps)	STS-3/STM-1	155.520 Mbps	35.699%
1000BASE-T (1000 Mbps)	STS-48/STM-16	2488.32 Mbps	59.812%

To summarize, the demand placed on the transport infrastructure by bandwidth-hungry applications and the explosive growth of the Internet has exceeded the limits of traditional TDM. Fiber, which once promised seemingly unlimited bandwidth, is being exhausted, and the expense, complexity, and scalability limitations of the SONET infrastructure are becoming increasingly problematic.

Wavelength Division Multiplexing

WDM increases the carrying capacity of the physical medium (fiber) using a completely different method from TDM. WDM assigns incoming optical signals to specific frequencies of light (wavelengths, or lambdas) within a certain frequency band. This multiplexing closely resembles the way radio stations broadcast on different wavelengths without interfering with each other (see Figure 1-7). Because each channel is transmitted at a different frequency, we can select from them using a tuner. Another way to think about WDM is that each channel is a different color of light; several channels then make up a "rainbow."

Figure 1-7 Increasing Capacity with WDM

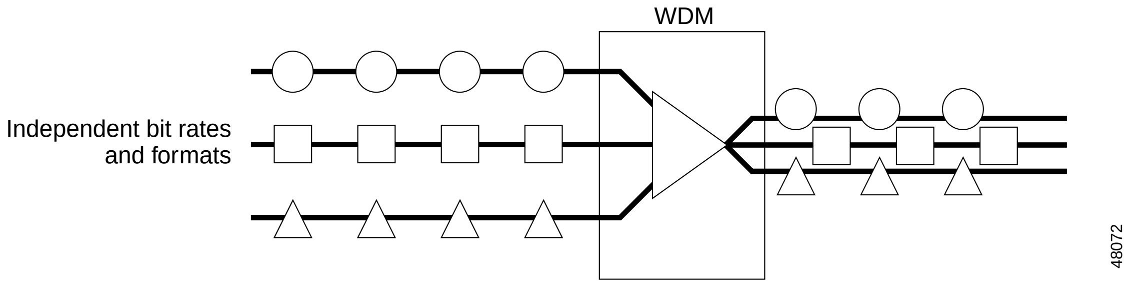

Note The term *wavelength* is used instead of the term frequency to avoid confusion with other uses of frequency. Wavelength is often used interchangeably with *lambda* and *channel*.

In a WDM system, each of the wavelengths is launched into the fiber, and the signals are demultiplexed at the receiving end. Like TDM, the resulting capacity is an aggregate of the input signals, but WDM carries each input signal independently of the others. This means that each channel has its own dedicated bandwidth; all signals arrive at the same time, rather than being broken up and carried in time slots.

The difference between WDM and dense wavelength division multiplexing (DWDM) is fundamentally one of only degree. DWDM spaces the wavelengths more closely than does WDM, and therefore has a greater overall capacity.The limits of this spacing are not precisely known, and have probably not been reached, though systems are available in mid-year 2000 with a capacity of 128 lambdas on one fiber. DWDM has a number of other notable features, which are discussed in greater detail in the following chapters. These include the ability to amplify all the wavelengths at once without first converting them to electrical signals, and the ability to carry signals of different speeds and types simultaneously and transparently over the fiber (protocol and bit rate independence).

Note WDM and DWDM use single-mode fiber to carry multiple lightwaves of differing frequencies. This should not be confused with transmission over multimode fiber, in which light is launched into the fiber at different angles, resulting in different “modes” of light. A single wavelength is used in multimode transmission.

TDM and WDM Compared

SONET TDM takes synchronous and asynchronous signals and multiplexes them to a single higher bit rate for transmission at a single wavelength over fiber. Source signals may have to be converted from electrical to optical, or from optical to electrical and back to optical before being multiplexed. WDM takes multiple optical signals, maps them to individual wavelengths, and multiplexes the wavelengths over a single fiber. Another fundamental difference between the two technologies is that WDM can carry multiple protocols without a common signal format, while SONET cannot. Some of the key differences between TDM and WDM are graphically illustrated in Figure 1-8.

Figure 1-8 TDM and WDM Interfaces

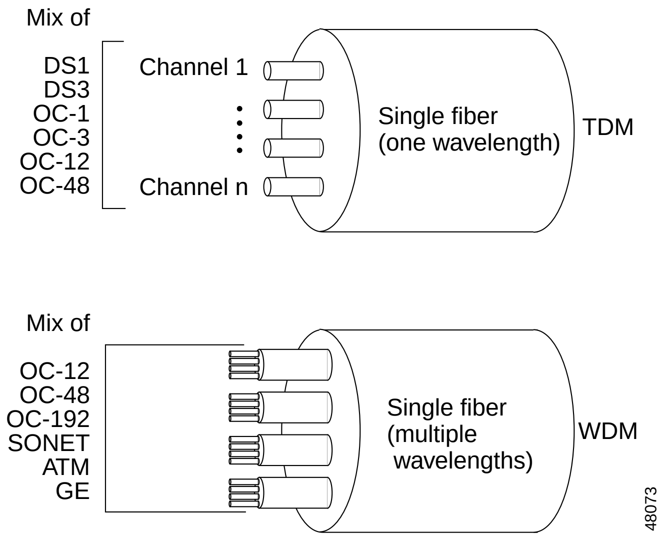

Additional Drivers in Metropolitan Area Networks

Bandwidth, the chief driver in the long-haul market, is also a big driver in metropolitan area, access, and large enterprise networks (see Figure 1-9). In these types of networks additional applications driving demand for bandwidth include storage area networks (SANs), which make possible the serverless office, consolidation of data centers, and real-time transaction processing backup.

Figure 1-9 High-Speed Enterprise WAN Bandwidth Migration

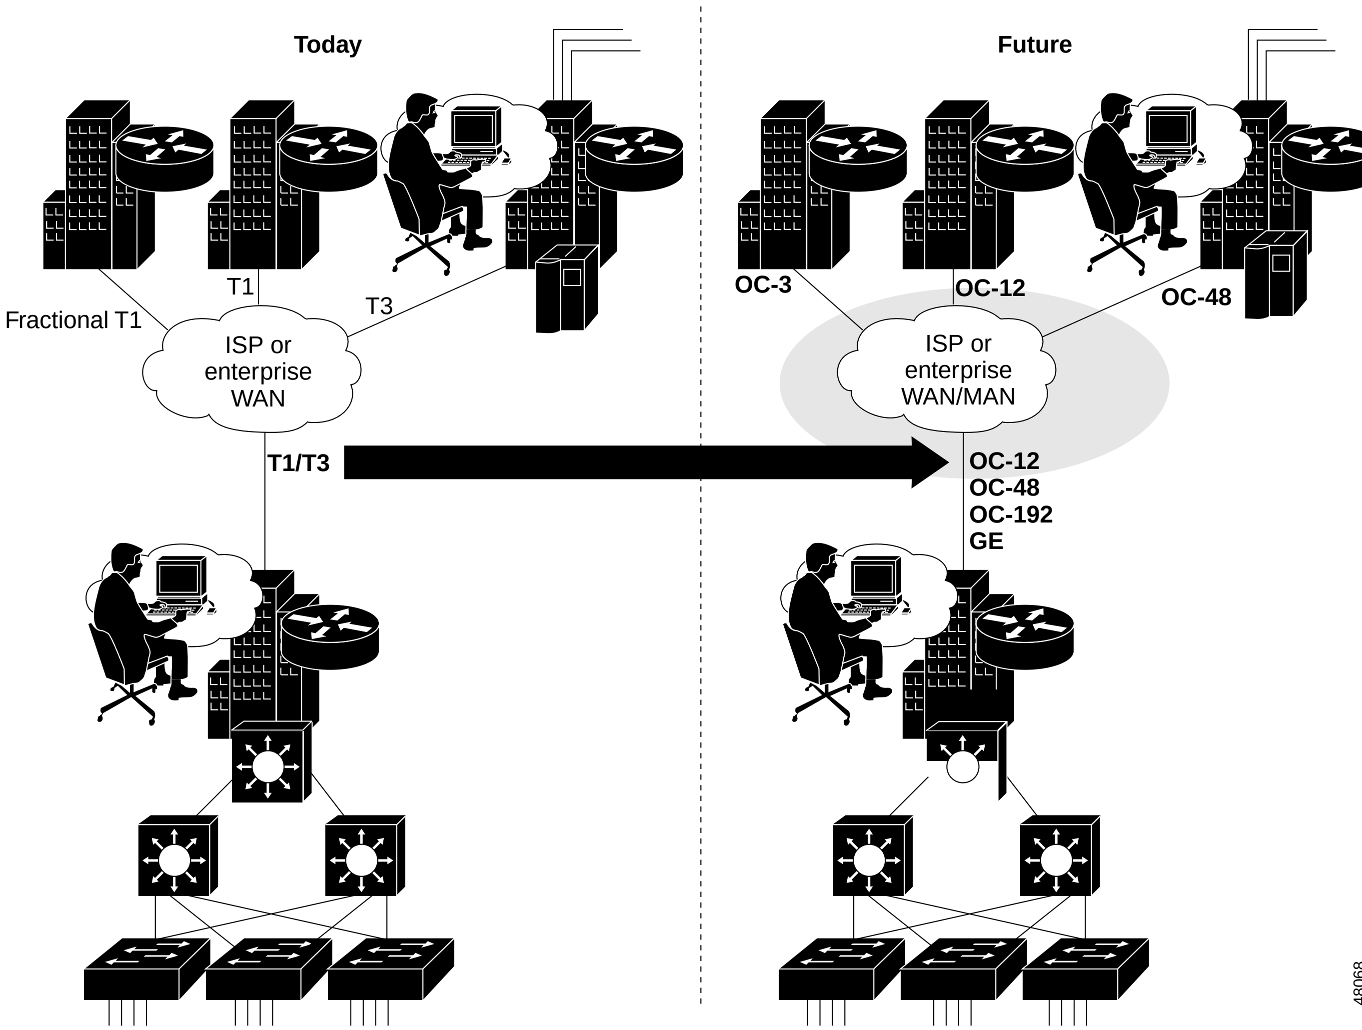

There is also rapidly increasing demand on access networks, which function primarily to connect end users over low-speed connections, such as dial-up lines, DSL, cable, and wireless, to a local POP. These connections are typically aggregated and carried over a SONET ring, which at some point attaches to a local POP that serves as an Internet gateway for long hauls. Now, the growing demand for high-speed services is prompting service providers to transform the POP into a dynamic service-delivery center. As a result, it is increasingly likely that a customer now obtains many high-speed services directly from the POP, without ever using the core segment of the Internet.

Value of DWDM in the Metropolitan Area

DWDM is the clear winner in the backbone. It was first deployed on long-haul routes in a time of fiber scarcity. Then the equipment savings made it the solution of choice for new long-haul routes, even when ample fiber was available. While DWDM can relieve fiber exhaust in the metropolitan area, its value in this market extends beyond this single advantage. Alternatives for capacity enhancement exist, such as pulling new cable and SONET overlays, but DWDM can do more. What delivers additional value in the metropolitan market is DWDM’s fast and flexible provisioning of protocol- and bit rate-transparent, data-centric, protected services, along with the ability to offer new and higher-speed services at less cost.

The need to provision services of varying types in a rapid and efficient manner in response to the changing demands of customers is a distinguishing characteristic of the metropolitan networks. With SONET, which is the foundation of the vast majority of existing MANs, service provisioning is a lengthy and complex process. Network planning and analysis, ADM provisioning, Digital Crossconnect System (DCS) reconfiguration, path and circuit verification, and service creation can take several weeks. By contrast, with DWDM equipment in place provisioning new service can be as simple as turning on another lightwave in an existing fiber pair.

Potential providers of DWDM-based services in metropolitan areas, where abundant fiber plant already exists or is being built, include incumbent local exchange carriers (ILECs), competitive local exchange carriers (CLECs), inter-exchange carriers (IXCs), Internet service providers (ISPs), cable companies, private network operators, and utility companies. Such carriers can often offer new services for less cost than older ones. Much of the cost savings is due to reducing unnecessary layers of equipment, which also lowers operational costs and simplifies the network architecture.

Carriers can create revenue today by providing protocol-transparent, high-speed LAN and SAN services to large organizations, as well as a mixture of lower-speed services (Token Ring, FDDI, Ethernet) to smaller organizations. In implementing an optical network, they are ensuring that they can play in the competitive field of the future.

Requirements in the Metropolitan Area

The requirements in the metropolitan market may differ in some respects from those in the long-haul network market, yet metropolitan networks are still just a geographically distinguished segment of the global network. What happens in the core must be supported right to the edge. IP, for example, is the dominant traffic type, so interworking with this layer is a requirement, while not ignoring other traffic (TDM). Network management is now of primary concern, and protection schemes that ensure high availability are a given.

Key requirements for DWDM systems in the MAN include the following:

- Multiprotocol support
- Scalability
- Reliability and availability
- Openness (interfaces, network management, standard fiber types, electromagnetic compatibility)
- Ease of installation and management
- Size and power consumption
- Cost effectiveness

Why DWDM?

From both technical and economic perspectives, the ability to provide potentially unlimited transmission capacity is the most obvious advantage of DWDM technology. The current investment in fiber plant can not only be preserved, but optimized by a factor of at least 32. As demands change, more capacity can be added, either by simple equipment upgrades or by increasing the number of lambdas on the fiber, without expensive upgrades. Capacity can be obtained for the cost of the equipment, and existing fiber plant investment is retained.

Bandwidth aside, DWDM's most compelling technical advantages can be summarized as follows:

- Transparency—Because DWDM is a physical layer architecture, it can transparently support both TDM and data formats such as ATM, Gigabit Ethernet, ESCON, and Fibre Channel with open interfaces over a common physical layer.
- Scalability—DWDM can leverage the abundance of dark fiber in many metropolitan area and enterprise networks to quickly meet demand for capacity on point-to-point links and on spans of existing SONET/SDH rings.
- Dynamic provisioning—Fast, simple, and dynamic provisioning of network connections give providers the ability to provide high-bandwidth services in days rather than months.

In the following sections we discuss some additional advantages, including migration from SONET and reliability.

SONET with DWDM

By using DWDM as a transport for TDM, existing SONET equipment investments can be preserved. Often new implementations can eliminate layers of equipment. For example, SONET multiplexing equipment can be avoided altogether by interfacing directly to DWDM equipment from ATM and packet switches, where OC-48 interfaces are common (see Figure 1-10). Additionally, upgrades do not have to conform to specific bit rate interfaces, as with SONET, where aggregation of tributaries is locked into specific values.

Figure 1-10 Direct SONET Interfaces from Switch to DWDM

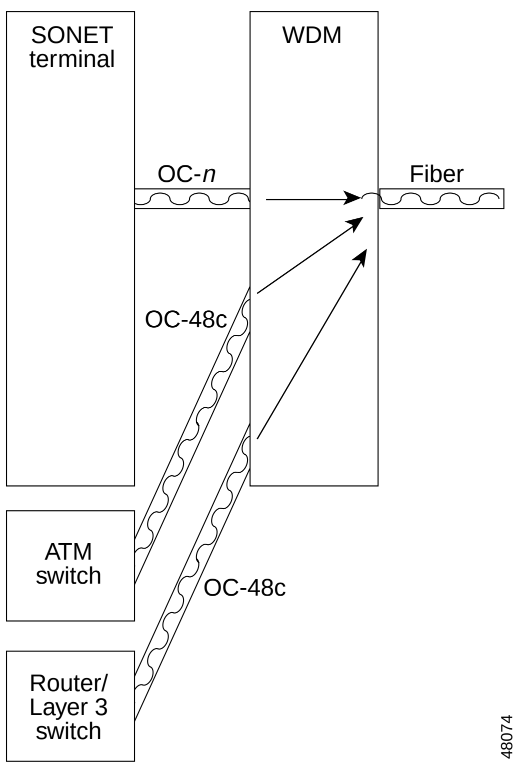

Optical signals become attenuated as they travel through fiber and must be periodically regenerated in core networks. In SONET/SDH optical networks prior to the introduction of DWDM, each separate fiber carrying a single optical signal, typically at 2.5 Gbps, required a separate electrical regenerator every 60 to 100 km (37 to 62 mi). As additional fibers were “turned up” in a core network, the total cost of regenerators could become very large, because not only the cost of the regenerators themselves, but also the facilities to house and power them, had to be considered. The need to add regenerators also increased the time required to light new fibers

The upper part of Figure 1-11 shows the infrastructure required to transmit at 10 Gbps (4 x OC-48 SR interfaces) across a span of 360 km (223 mi) using SONET equipment; the lower part of the figure shows the infrastructure required for the same capacity using DWDM. While optical amplifiers could be used in the SONET case to extend the distance of spans before having to boost signal power, there would still need to be an amplifier for each fiber. Because with DWDM all four signals can be transported on a single fiber pair (versus four), fewer pieces of equipment are required. Eliminating the expense of regenerators (RPTR) required for each fiber results in considerable savings.

Figure 1-11 DWDM Eliminates Regenerators

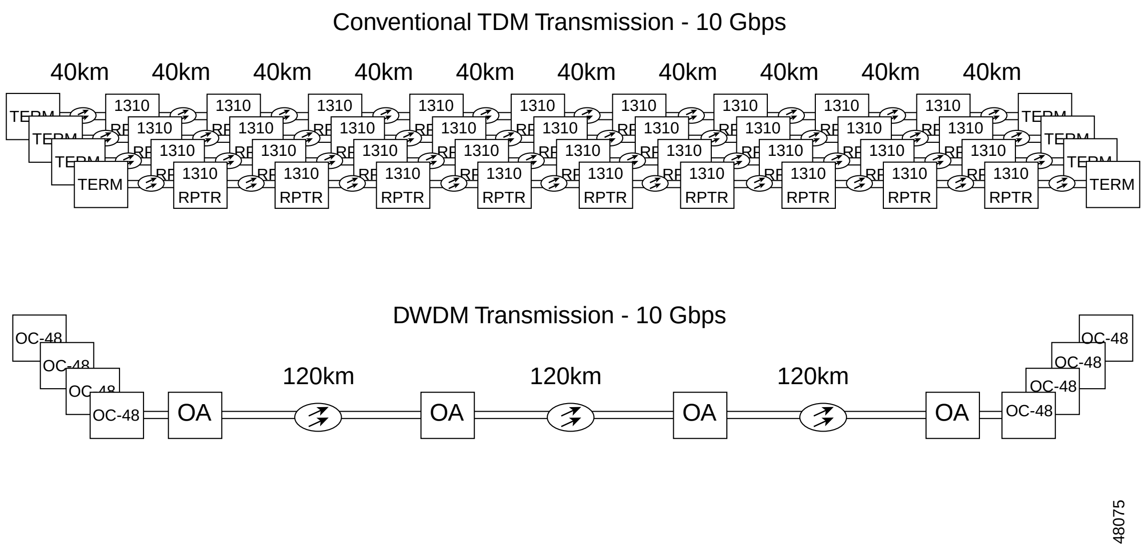

A single optical amplifier can reamplify all the channels on a DWDM fiber without demultiplexing and processing them individually, with a cost approaching that of a single regenerator. The optical amplifier merely amplifies the signals; it does not reshape, retime or retransmit them as a regenerator does, so the signals may still need to be regenerated periodically. But depending on system design, signals can now be transmitted anywhere from 600 to thousands of kilometers without regeneration.

In addition to dramatically reducing the cost of regenerators, DWDM systems greatly simplify the expansion of network capacity. The only requirement is to install additional or higher bit-rate interfaces in the DWDM systems at either end of the fiber. In some cases it will only be necessary to increase the number of lambdas on the fiber by deploying existing interfaces, as shown in the upper half of Figure 1-12. The existing optical amplifiers amplify the new channel without additional regenerators. In the case of adding higher bit-rate interfaces, as shown in the lower half of Figure 1-12, fiber type can become a consideration. See the “Optical Fibers” section on page 2-5 for an overview of types of optical fibers and their uses.

Figure 1-12 Upgrading with DWDM

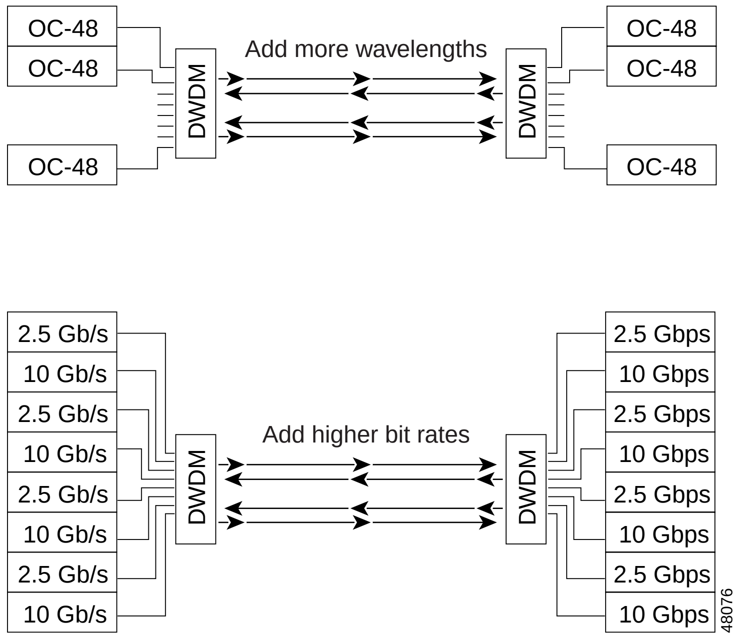

Although amplifiers are of great benefit in long-haul transport, they are often unnecessary in metropolitan networks. Where distances between network elements are relatively short, signal strength and integrity can be adequate without amplification. But with MANs expanding in deeper into long-haul reaches, amplifiers will become useful.

Enhancing Performance and Reliability

Today's metropolitan and enterprise networks support many mission-critical applications that require high availability, such as billing and accounting on mainframes or client-server installations in data centers. Continuous backups or reliable decentralized data processing and storage are essential. These applications, along with disaster recovery and parallel processing, have high requirements for performance and reliability. As enterprises out source data services and inter-LAN connectivity, the burden of service falls on the service provider rather than on the enterprise.

With DWDM, the transport network is theoretically unconstrained by the speed of available electronics. There is no need for optical-electrical-optical (OEO) conversion when using optical amplifiers, rather than regenerators, on the physical link. Although not yet prevalent, direct optical interfaces to DWDM equipment can also eliminate the need for an OEO function.

While optical amplifiers are a major factor in the ability to extend the effective range of DWDM, other factors also come into play. For example, DWDM is subject to dispersion and nonlinear effects. These effects are further discussed in the "Optical Fibers" section on page 2-5.

Many components, such as the optical add/drop multiplexer (OADM), are passive and therefore continue to work, even if there is a power cut. In addition, these components tend to have a very high mean time between failures (MTBF). Protection schemes implemented on DWDM equipment and in the network designs are at least as robust as those built into SONET. All these factors contribute to better performance and lower maintenance in the optical network.

Network Management Capability

One of the primary advantages offered by SONET technology is the capability of the data communication channel (DCC). Used for operations functions, DCCs ship such things as alarms, administration data, signal control information, and maintenance messages. When SONET is transported over DWDM, DCCs continue to perform these functions between SONET network elements. In addition, a DWDM system can have its own management channel for the optical layer. For out-of-band management, an additional wavelength (for example, a 33rd wavelength in a 32-wavelength system) can be used as an optical supervisory channel. For inband management, a small amount of bandwidth (for example, 8 kHz) can be reserved for management on a per-channel basis.

Additional Benefits

The shift in the makeup of traffic from voice to data has important implications for the design and operation of carrier networks. The introduction of cell-switching technologies such as ATM and Frame Relay demonstrates the limitations of the narrow-band, circuit-switched network design, but the limits of these technologies are being reached. Data is no longer an add-on to the voice-centric network, but is central. There are fundamentally different requirements of a data-centric network; two of these are the aggregation model and the open versus proprietary interfaces.

Aggregation in a voice-centric network consists of multiplexing numerous times onto transmission facilities and at many points in the network. Aggregation in a data-centric network, by contrast, tends to happen at the edge. With OC-48 (and higher) interfaces readily available on cell and packet switches, it becomes possible to eliminate costly SONET multiplexing and digital cross-connect equipment. OC-48 connections can interface directly to DWDM equipment.

Finally, service providers and enterprises can respond more quickly to changing demands by allocating bandwidth on demand. The ability to provision services rapidly by providing wavelength on demand creates new revenue opportunities such as wavelength leasing (an alternative to leasing of physical links or bit rate-limited tunnels), disaster recovery, and optical VPNs.