

Brought to you by **ciena**

Imtech

Imtech Telecom Global

Optical Networking

FOR
DUMMIES[®]

**A Reference
for the
Rest of Us!**[®]

FREE eTips at dummies.com[®]

Get the ins and
outs of optical
networking



Ed Tittel
with Chris Janson

Optical Networking FOR **DUMMIES®**

by Ed Tittel, with Chris Janson



WILEY

Wiley Publishing, Inc.

Optical Networking For Dummies®

Published by

Wiley Publishing, Inc.

111 River Street

Hoboken, NJ 07030-5774

Copyright © 2009 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For details on how to create a custom *For Dummies* book for your business or organization, contact bizdev@wiley.com. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-0-470-44759-8

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



WILEY

Table of Contents

Introduction	1
Conventions Used in This Book	1
Icons Used in This Book	1
How This Book Is Organized	2
Chapter 1: What Is Networking?	3
What Is a Network?	3
What Makes a Network Work?	4
When Computers Talk to Each Other, What Do They Say?	5
Stacking Up to Networking	6
A Brief Overview of Networking	11
How Is a Modern Network Built?	12
How Do Modern Networks Operate?	14
How Network Management Really Works	15
Chapter 2: Using Light to Communicate	17
What Is an Optical Network?	18
Fiber Optics 101	25
Chapter 3: Building Bigger Networks.	31
Double, Triple, and Tenfold Bandwidth	31
Another Way of Multiplexing	34
Core Terminologies and Principles	40
Doing More with the Same Fiber	43
Chapter 4: Making the Most of Your Optical Network	45
Grooming Is More than Good Hygiene	45
Traffic Grooming Techniques	48
Is This Thing Reliable?	50
Optical Network Topologies	52
Chapter 5: Ten (Or More) Views of the Future	57
Telco and Data Networks: Strange Bedfellows	58
A View of the Unified Network	63
Ten Prognostications on Optical Networking	65
A Glossary of Acronyms	67

Publisher's Acknowledgments

We're proud of this book; please send us your comments through our Dummies online registration form located at <http://dummies.custhelp.com>. For other comments, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For details on how to create a custom *For Dummies* book for your business or organization, contact bizdev@wiley.com. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Media Development

Senior Project Editor: Zoë Wykes

Editorial Manager: Rev Mengle

Business Development Representative:
Sue Blessing

Custom Publishing Project Specialist:
Michael Sullivan

Production

Senior Project Coordinator: Kristie Rees

Layout and Graphics: Reuben W. Davis,
Melissa K. Jester

Proofreaders: Melissa Cossell

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher

Andy Cummings, Vice President and Publisher

Mary Bednarek, Executive Director, Acquisitions

Mary C. Corder, Editorial Director

Publishing and Editorial for Consumer Dummies

Diane Graves Steele, Vice President and Publisher, Consumer Dummies

Composition Services

Gerry Fahey, Vice President of Production Services

Debbie Stailey, Director of Composition Services

Introduction

If you've ever wondered how light can ferry vast amounts of information across huge distances, wonder no more. This book is just for you. Here, you discover how science and technology enable modern networks to convey humongous hunks of information (more than the entire contents of the Library of Congress, every second) hither, thither, and yon.

Though optical networking may sound esoteric, it really isn't. Optical networking relies on precision engineering and the ability to do lots of things very, very fast. That's where the speed of light works in its favor. In this book, you find out how optical networks exploit basic physical properties of light energy to make networking reliable at high speeds and tremendous volumes. Welcome to *Optical Networking For Dummies*!

Conventions Used in This Book

Italic type indicates a defined term or something we want to emphasize. **Boldface** highlights the key word in a bulleted list and the action in a set of steps.

You see text in gray boxes occasionally. These *sidebars* contain information you may find interesting, but you don't need to master it to understand the topic at hand. So, you can skip the boxes altogether, go back and read them at your leisure, or read them along with the regular text.

Icons Used in This Book

Every *For Dummies* book has small illustrations, called *icons*, sprinkled throughout the margins. These tiny images call attention to text we think is worth special attention for one reason or another. Following are the icons we use in this book.



Points to keep in mind as you immerse yourself in the world of optical networking are highlighted with this icon.



We're networking aficionados, and we really like sharing our knowledge with you. But, we realize that you don't necessarily need to know everything we do, so this icon tells you that the text here is more detailed than is absolutely necessary, and you can skip it if you like.



Right-on-target information you can use to help you make the most of any investment in optical networking appears beside this bull's-eye.

How This Book Is Organized

The five chapters in this book lead you into the components and technologies used in optical networking, and then some:

- ✓ **Chapter 1, What Is a Network?**, explains basic network capabilities, and how networks are built and operated.
- ✓ **Chapter 2, Using Light to Communicate**, looks into the optical networking technologies and hardware.
- ✓ **Chapter 3, Building Bigger Networks**, explores how increasing network-carrying capacity boosts capabilities.
- ✓ **Chapter 4, Making the Most of Your Optical Network**, explains how to manage and monitor optical networks most effectively.
- ✓ **Chapter 5, Ten (Or More) Views of the Future**, tells you how voice, video, and data converge into a unified network.
- ✓ **A Glossary of Acronyms**, helps you to convert arcane abbreviations into something you can understand.

The chapters are designed to stand alone, so if you're dying to know about monitoring a network, head straight to Chapter 4; if you want to see the light, turn to Chapter 2; or just turn the page and keep on going.

Chapter 1

What Is Networking?

In This Chapter

- ▶ Looking at a network
 - ▶ Making networks work
 - ▶ Touring a typical network
 - ▶ Operating a network
-

If you've ever built a tin-can-and-string telephone, you already know more about networking than you might think. Although computer networks may seem infinitely more mysterious than "Mr. Watson, come here. I need you" (Alexander Graham Bell's first words on the telephone to his assistant), sooner or later all networks boil down to providing a communications link between computers that works as it should.

This chapter introduces networking, a fundamental communications capability and service that enables modern working life (and a lot of play) as we know it.



If you think tin-can-and-string telephones have nothing in common with networking, remember that any phone depends on somebody talking while somebody else listens. Yelling is optional, but sometimes helps!

What Is a Network?

In more formal terms, a *network* is a collection of computers that are linked together so they can communicate. Most networks use some kind of cable — typically copper wire but fiber optic and even wireless varieties are also available — to link computers and other devices through a connection

that permits the computer to listen and talk to the wire. More than just hardware is involved in this exchange. Cables and connections are essential to networking, but without software they're purely decorative.

Networking imposes three fundamental requirements to do its job of making devices communicate:

- ✓ **Connections** include the physical bits of gear needed to hook a computer to a network and the wires or other materials — known as the networking medium — that carry messages from one computer to another. Without physical connections, computers are isolated from the network and have no way to interact.
- ✓ **Communications** establish the rules for how computers talk to each other and what things mean. Because one computer can run radically different software than another one does, computers must speak some kind of “shared language” to enable them to talk to each other. Without shared communications, computers can't exchange information with each other, and they remain isolated.
- ✓ **Services** define the things that computers can talk about with each other and what they can *do* for each other, such as sending or receiving files, exchanging messages, looking up information, and so on.

What Makes a Network Work?

The three networking fundamentals — connections, communications, and services — must come together for a network to function. But what really makes a network work?

First, the physical connections have to work so any computer can talk or listen to the network medium that serves as the highway over which signals can move between computers. Without a working connection, no signals can enter or leave a computer or other device.

Second, the communications have to work so that when one computer talks to another, the sender knows what signals to emit — and the receiver(s) know how to listen to

and interpret incoming signals. Without working communications, nothing meaningful happens because there's no way to make sense of incoming signals or to create outgoing ones.

Third, computers must be able to work together so that one can ask for things that another can deliver, and vice versa. Without a shared set of services, neither computer can do anything for the other — and nothing happens.

When Computers Talk to Each Other, What Do They Say?

When computers communicate with each other, they generally spend a small amount of time at the outset establishing a connection. After that, they spend the bulk of their time exchanging information about a specific service, often consisting of a series of one or more requests and replies between the parties involved. Once that exchange comes to an end, they spend another small snippet of time breaking their connection so others can use the resources they've been consuming while their connection was active.

Do what I mean, not what I say!

Computers can only do exactly what they're told to do, neither more nor less. For computers to communicate, and for any information to be delivered, every bit of information that the sender and receiver want to exchange must be supplied explicitly.

For computers to communicate, they must share a set of rules that lets them do all of the following:

- ✓ Obtain and use a unique network address for self-reference and identification
- ✓ Understand how to obtain and use other unique addresses to communicate with other devices
- ✓ Understand how to grab and use the network medium for communication (this means knowing how and when to take a turn, how long to talk, when to stop talking, how to send and receive busy signals, and how to wait in line for a turn to talk)

Defining the rules means you must completely map out these issues in software to formulate a complete set of behaviors to let computers communicate. Such sets of rules are called *networking protocols*, and they operate at various levels to permit computers to access and use networks to send and receive messages of all kinds.

Setting the standards

In the early days of networking, equipment and software vendors made up the rules for networking as they went along. As long as you bought everything you used for networking from Vendor A or Company B, all was okay. But, as soon as buyers tried to start mixing and matching offerings from multiple providers, they quickly realized that incompatibilities prevented their gear from working together as they wanted.

If this story had a truly happy ending, it might go something like, “Today, there’s only one set of networking rules.” Alas, that’s not quite true. The chaos has been reduced, but there are enough standards around to leave room for lots of occasional confusion.

Stacking Up to Networking

The age of networking begins with work undertaken as far back as the early 1970s by that era’s computing giants, such as IBM and Sperry. In addition to these players — and a powerful cadre of researchers and developers working for the government and companies like Xerox and others — networking continued to explode in the late 70s and 80s as a growing group of smaller players (Digital Equipment Corporation, also known as DEC, Data General, 3COM, and others) got involved.

In many ways, the early evolution of networking parallels the early days of the automobile. In both cases, these markets were substantially fragmented at the outset, with hundreds of small, innovative companies competing to sell products as fast as they could. But, as consumers jumped on these offerings, all players soon realized that both networking and automobiles require substantial amounts of infrastructure to reach everywhere we want them to go.

What followed next was substantial consolidation in both markets. For autos, it meant that hundreds of companies became the “Big Three” within 30 years. For networking, it meant that networking standards drove all players who survived into a small and relatively manageable group of technologies, each defined by a shared networking standard for some particular market niche, communications technology, or related networking protocols and software.

A protocol's work is never done

Okay, so now you know that one computer cannot talk to another without sharing a common protocol. But where does this protocol stuff come from? Or, in different terms, who decides which protocols to use?

In fact, protocols span the range of networking all the way from software to hardware. The programs that permit your computer to access the network *must use* a protocol to interact with another network. This protocol continues all the way down to the very edge of the hardware, where the computer says “send this series of signals” to talk to the network or “give me the signals” when the hardware informs the computer that incoming data has arrived from the network.

Onions have layers, networks have layers

Network protocols are like onions in that they consist of a series of separate layers arranged in a logical hierarchy. In the aggregate, networking technology qualifies as a “Big Problem,” because there are lots of tasks and activities involved in creating and using a working network. When faced with such problems, scientists and engineers typically employ a technique called “Divide and Conquer” to subdue these otherwise savage and intractable beasts.

For networking, divide and conquer has had differing historical interpretations and numerous implementations, but all are founded on this understanding: The engineers and technologists who understand and build the best networking hardware

(what some people like to call the “plumbing” aspect of networking) are not the same engineers and technologists who understand and build the software that enables computers to use that hardware to provide services, exchange messages, complete transactions, and all the other stuff that networks actually do.

Networking technologies have been broken up into between four and seven layers to put “divide and conquer” effectively to work. The most popular four-layer model in use today comes from the earliest days of networking as practiced by IBM and various U.S. government projects (still used today in the very set of higher-level protocols that make the Internet work). It’s usually called the *DoD network reference model*.

The reigning seven-layer model is one of the only surviving and still thriving results of a massive international standards effort called the *Organization for International Standards Open Systems Interconnection Initiative*. This effort is usually identified as OSI/ISO. The related network model is usually called the ISO/OSI reference model, or the OSI stack.

Figure 1-1 shows how these two models stack up against each other visually (we explain their layers following the figure):

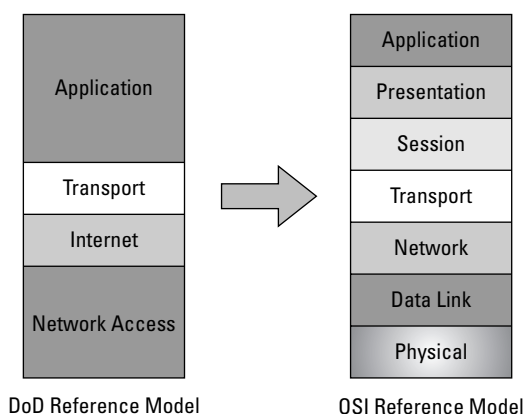


Figure 1-1: The original DoD reference model and today's OSI model.

Here's another simplification worth considering before we start decoding layers. The bottom layers (Data Link and

Physical for the OSI Reference Model, Network Access for the DoD model) focus on interacting with networking hardware, managing access to the networking medium, and sending and receiving signals across that medium. The top layers (Session through Application in the OSI Reference Model) concern themselves with interacting with computer software, managing information transfer onto and off the network.

Thus, it's important to understand that both models talk to local computer or device software (applications) above the stack, and talk to local network hardware devices (interfaces) below the stack. Everything else is protocol, and that's where various layers come into play.

The way the two diagrams line up is deliberate. This means that what happens in either model is roughly the same, give or take a few details, at the same level. This also explains why the OSI model is the most widely used tool for explaining networking today, even though the DoD model is actually more widely deployed. The OSI model helps add clarity and detail and makes for a good explanation, even if the DoD model describes how the Internet protocol suite actually works. These DoD protocols are best known as TCP/IP, for *Transmission Control Protocol/Internet Protocol*.

OSI Physical layer (1)

The *Physical* layer 1 includes the physical network medium (copper or optical fiber cables, wireless media) that any network uses to send and receive signals across the medium. It also embraces the details involved in such signaling, and the physical and electrical (and related optical or radio frequency) characteristics of the network interfaces that connect computers and other devices to the network medium, and vice versa. A quick simple summary might be "it's about hardware."

OSI Data Link layer (2)

This layer's job is to facilitate reliable transmission of data through the Physical layer when sending, and to check reliability of signals when receiving. The Data Link layer also handles point-to-point network transmission across the medium from one sender to a single receiver. Basically, the *Data Link* layer translates bits into signals for outgoing messages, and vice versa for incoming ones. It also handles media access control (MAC) and related hardware addresses,

as well as logical link control (LLC) between pairs of communicating network interfaces. A quick summary is “converts signals into bits, and bits into signals.”

Network layer (3)

This is where network location information and the traffic control involved in moving or routing messages from sender to receiver are handled. The *Network* layer handles logical network addresses associated with devices or computers to correlate human-readable names for them with binary machine-readable addresses. This same addressing information is also used to determine how messages should travel from sender to receiver if all parties do not reside on the same cable segment (in that case, LLC handles this at the Data Link layer). Finally, the Network layer can identify which processes on a computer or device are communicating with the network, and direct traffic accordingly — which lets users conduct multiple simultaneous network sessions (e-mail plus one or more Web browser windows or tabs, for example). Up to this layer, message traffic is optimized for best network traversal, too. The quick summary is “network location and routing from sender to receiver.”

Layers 4 through 7: Transport to Application

These layers are responsible for allowing reliable communications between senders and receivers, and for coordinating ongoing information exchange (called a *session*) between communication partners. These layers also make sure that what one computer sends out is intelligible to the computer on the receiving end, and help computers to map network services to related applications and vice versa. As it happens, none of these things is terribly important to optical networking — though of course, they’re terribly important to actual network users — but hopefully that explains why we gloss over them so quickly here.

Relating the OSI and DoD network models

It’s absolutely fair to equate the functions ascribed to the OSI layers on the left side of the preceding chart to the DoD layers on the right side of that chart. Thus, Data Link and Physical (OSI) correspond to Network Access (DoD), Network (OSI) corresponds to Internet (DoD), Transport means the same to

both models, and Application, Presentation, and Session layers (OSI) map into a single more monolithic Application layer (DoD).

Otherwise, the capabilities and functions are similar enough from one model to the next that we won't belabor the point by rehashing the functions already described for the OSI model once again for its DoD counterpart.

A Brief Overview of Networking

The earliest networks were introduced in the 1960s, strictly in a vendor specific (and often per-platform) format. By the early 1970s, some standardization had occurred for telephone connections, as devices called *modulators/demodulators* (better known as modems) were introduced to let users establish sessions with mainframe computers.

By the end of the 1970s, engineers were developing standards for various types of networking hardware that included Ethernet (still in use today in many forms), Token-Ring, and ARCnet (now ancient history). At the same time, long-haul networks advanced from single telephone connections to incorporate various forms of trunk lines. This is when the distinction between local-area networking or LAN (what companies and organizations do on their premises) and wide-area networking or WAN (what companies and organizations do when they must span long distances, generally also involving a phone company or telecommunications carrier to provide the intervening "long distance" infrastructure) technologies emerged.

Since the 1980s, networking shows three major trends:

- ✓ **Ever-increasing LAN speeds.** For example, the earliest version of Ethernet ran at 3 megabits per second (Mbps) at Xerox in 1975, which turned into 10 Mbps as coaxial versions appeared in 1980 (an early and successful example of standards-based networking). 100 Mbps Ethernet followed in the early 90s, and 1000 Mbps Ethernet (*also known as* Gigabit Ethernet or GbE) followed in the late 90s. 10 GbE is taking hold in 2008 while 40GbE and 100GbE are emerging standards.

✓ **Ever-increasing WAN speeds.** Early telephone links jumped from 300 to 1200 to 2400 baud (a measure close enough to bits per second or bps that we don't distinguish them here, though there are technical differences) in the 70s and 80s. Trunk lines at 1.544 Mbps started to become popular in the mid-80s (also known as T-1 and DS1 in the U.S., with a 2.048 formulation known as E-1 popular in Europe and elsewhere around the world). In the 90s, long haul links at DS3 levels (28 DS1, or about 45 Mbps) first became available for WAN networking. In the period from the 70s to the 90s, a nascent optical networking infrastructure (more on this in Chapter 2) began to emerge with optical carrier data rates starting at OC-1 (about 52 Mbps) up to OC-24 (about 1.2 Gbps). Today, OC-levels from 48 (2.4 Gbps) to OC-768 (about 39 Gbps) make up the backbone of the carrier infrastructure. Moving forward, 100Gbps will emerge in carrier and WAN backbones.



As of December 2007, AT&T alone claimed 50,000 wavelength-miles of OC-768 in its Internet/MPLS backbone network, and OC-768 interfaces have been available from network vendors since 2006.

✓ **Increasing mainstreaming of networking technologies.** Especially for LAN Ethernet, even GbE and for various wireless networking technologies, networking is increasingly built in on computing equipment. Most notebook PCs ship with wireless and wired Ethernet connections, desktop PCs with one or two GbE connections, and server PCs with two or more GbE connections.

Networking is everywhere that PCs are nowadays, and it's fast and cheap. This is what drives the appetite for ever more Internet access, and opens the door for "rich applications" that enable voice and video communications for PCs. In fact, at the very highest speeds most development for both LAN and WAN technologies focuses on 100 GbE efforts underway through the Institute of Electrical and Electronics Engineers (IEEE, pronounced eye-triple-E) Higher Speed Study Group.

How Is a Modern Network Built?

These days when companies or organizations construct networks, they build around serious infrastructure. The

wide-area/local-area (WAN/LAN) distinction is still important, but the “need for speed” on both sides of the network boundary that divides them drives the inclusion of serious and powerful elements all along the way.

Building modern LANs

Inside the Internet boundary — positioned by convention at the “edge” of an organization, right between the LAN and WAN worlds, modern organizations use devices called *routers*. These devices basically look at network traffic and send it where it needs to go, and sit between the Internet or WAN side and the LAN side of the network to manage outgoing and incoming traffic. Routers are typically large powerful devices that can manage hundreds of thousands to millions of connections at once, and also provide various levels of security, both on their own and in tandem with special *firewalls* that secure the network boundary.

At the core of the enterprise or organizational LAN sits one or more special, high-speed networking devices called *switches*. These devices can handle multiple simultaneous connections between pairs of computers or other devices, generally from thousands to tens of thousands of connections at any given moment. This is usually sufficient for even the largest local-area network environments.

Other aspects of enterprise or organization network infrastructure nowadays also include storage networks, which provide network-wide access to file services for everyone, and which also provide high-speed block level storage for servers on such networks. Increasingly, it’s becoming cost effective to centralize storage in the data center for all kinds of uses. Such storage networks often use their own special, high-speed networking infrastructure in parallel with conventional networks for access to other services. This approach also makes centralized backup easier at the large end of the scale, where conventional networks might otherwise be swamped by periodic backup traffic.

Building a modern WAN

Until recently, companies and organizations in need of WAN access had two basic options:

- ✓ To contract with one or more communications carriers to lease WAN capacity, as a way to create private site-to-site long-haul network links
- ✓ To contract with an Internet Service Provider (ISP) to obtain Internet access, sometimes in tandem with a communications carrier for a link from the customer premises to the ISP site

Both methods essentially involve up-front costs to put the link or links in place, then metered (pay-as-you-go) costs for “unlimited bandwidth” or fixed monthly costs for bandwidth and total utilization caps. Both methods also put the company or organization making a contract into a long-term relationship with the carrier and/or ISP, subject to the ups and downs of such relationships. They must rely on service level agreements and constant monitoring to make sure they get the service they pay for, and on the provider’s infrastructure to deliver the networking goods.

Today, another WAN option is open to adventurous, technically savvy organizations and companies. They can lease so-called “dark fiber” — fiber optic cable that is in place but not in use — from a carrier or telco, and build their own optical infrastructure. This means they have to buy the expensive WAN networking gear necessary to “light up” the dark fiber and operate their own related infrastructure. But this confers the ability to use as much bandwidth as that equipment and fiber optic infrastructure can deliver, without incurring extra costs. It also lets companies depreciate the costs of this infrastructure over time, rather than treating monthly service costs as a business expense.

How Do Modern Networks Operate?

Modern networks include various types of monitoring and management capabilities in all of the devices on the network, especially those devices (such as switches, routers, firewalls, and so forth) that define the network’s structure. Each individual component reports into some kind of network

management console to track errors, availability, uptime, utilization, and other key metrics to help monitor and report network health.

Of course, in environments where service level agreements (or SLAs) prevail, such metrics also establish when or if service levels are being met, and to measure deviations from what those agreements prescribe. Various other measurement tools may be used to track how end-users perceive applications to be working (often called “user experience measurement”) to make sure that those who actually use such networks can get their jobs done.

In larger networks, there may even be a hierarchy of management consoles. The consoles that provide big-picture information may be known as “manager of managers” (MOM) or end-to-end management tools.

How Network Management Really Works

At the highest level of network management, operators look for signs that things aren’t working as they should be, visually presented through an easily understood dashboard. There, operators will typically find read-outs that depict the following sorts of things:

- ✓ **Network availability:** Indicates whether the network is up and running.
- ✓ **Traffic information:** Shows the level of network load, the types of traffic involved, and deviations from expected/typical loads and mixes.
- ✓ **Service levels:** Show whether SLA requirements are being met (green), are on the edge (yellow), or have crossed various failure thresholds (red).
- ✓ **User experience:** Indicates whether typical interactions or real-time application measurements are acceptable (green), on the edge (yellow), or have fallen below various failure thresholds (red).

Ultimately this boils down to a dashboard display that might look something like what’s shown in Figure 1-2.




Figure 1-2: A typical dashboard zeroes in on availability and service-level measures, along with service-quality history over time and business calendar data.

Chapter 2

Using Light to Communicate

.....

In This Chapter

- ▶ Understanding optical networking technologies
 - ▶ Taking a turn down optical networking's memory lane
 - ▶ Introducing optical networking fundamentals and technologies
-

Although it may seem strange to use light as a means of communication, you actually do this every day. Think about it: When you flick your high beams at an oncoming driver to ask the person to turn down his lights, that flick itself is a form of communication. When you sit at a traffic signal, it's the switch from red to green that tells you it's time to go. When accidents or breakdowns occur at night, bright flares warn oncoming motorists of trouble ahead. Runway lights provide key positioning and orientation information to let pilots land planes in the dark. In short, light can deliver all kinds of useful and even essential life-or-death information.

Optical technologies also play a huge role in networking. In this case, most forms of such communication use special optical fibers as a communications medium to link devices that can emit light on one end of an optical fiber to send information and to detect light on the other end of the optical fiber to receive it. In this chapter, we dig into the fundamentals that make this technology work, explain what it's made of, and explore its possible applications.

What Is an Optical Network?

Of course, there's more to creating optical networks than simply stringing optical fibers from point to point, then shining bright lines into them to communicate. In keeping with our earlier definition of networking (see Chapter 1), we know that all networks:

- ✓ Require physical connections and devices to send and receive signals
- ✓ Use a set of shared rules that govern how signals are sent and received to establish and manage communications
- ✓ Define a set of services that employ physical connections and devices and rules for communication to create a working network that can actually *do* something

In the following sections, we explore what these things mean in the context of optical networking.

Optical network devices and media

Okay, now you know that optical communications use light to communicate. That said, optical networking involves three major components to support optical communications:

- ✓ Lasers that produce light at very precise wavelengths (these are the devices that emit light to send information)
- ✓ Optical fibers that transport light from one place to another (or, more precisely, from one end of the fiber to another)
- ✓ Receivers or sensors that detect this light

Of course, there's another uninvited guest who always shows up at this party, much like ants at a picnic — namely, dirt. Dirt is everywhere, and dirt blocks light, so working with fiber optic technologies and devices — especially mating connectors with the ends of fiber optic cables — takes extreme cleanliness. If the light is to go through as it must, no dirt can be allowed to get in the way!

Lasers emit light

As it turns out, the word *laser* is an acronym. It stands for Light Amplification by Stimulated Emission of Radiation (LASER). One basic laser type works by applying sizable amounts of energy to a quartz flash tube, which emits bright light that excites photons inside a crystal that is covered in a 95 percent reflective coating, with a near-perfect reflective coating on one end of the crystal (as close to 100 percent reflectivity as modern technology allows, in fact).

As photons bounce around inside the crystal, the intensity of the contained light increases through a process of amplification called *pumping*. Once that light energy reaches an appropriate level, or some specific time interval has passed (that allows enough energy to accumulate), the laser emits an intense beam of light through an output coupler on the other end of the crystal. The reflector isn't perfect either or it wouldn't be able to allow the light to escape. A primitive ruby laser is depicted in Figure 2-1.

Components of the first ruby laser




Figure 2-1: A ruby laser applies external energy to a ruby crystal that builds up photons until they emit as an intense beam of (red) light from the output coupler.



Typical optical networking lasers

Several types of lasers are widely used in modern optical networking hardware. Vertical Cavity Surface Emitting Lasers (VCSEL) are computer chips that emit light straight up. They're made of gallium arsenide and emit light in the 750-860 nm range. Fabry Perot or FP lasers produce only 1310 nm outputs, and produce lots of individual bandwidths in a narrow range.

Distributed Feedback (DFB) lasers are used for both 1310 and 1550 nm applications up to about 2.5 Gbps, produce tight wavelengths, and can

send light over longer distances. An externally modulated laser (one type is an electro-absorption modulated laser, or EML) is also semiconductor-based but much bigger than a VCSEL. These devices are normally cooled to maintain constant temperatures (or wavelengths of light output will vary), and are used for high-speed applications from 2.5 Gbps to 10 Gbps. Higher bit rate systems use a DFB laser and a separate external modulator, usually made of Lithium Niobate to reach rates up to 100Gbps. These usually operate at 1310 or 1550 nm wavelengths.

Though the earliest lasers used crystalline substances to accumulate photons for emission, modern lasers can use all kinds of materials to do this job. Generically called the *gain medium*, this part of a laser is where photons accumulate in response to energy pumping from some external energy source. Modern lasers may use gases, liquids, solids, or even high-energy plasmas as the gain medium.



In a laser, the gain medium is chosen because of its ability to absorb pump energy. This raises electrons in the medium into a higher-energy quantum state. Such particles interact with light by emitting photons through direct stimulation from the pump. This causes the amount of light to increase, or be amplified. By placing the gain medium inside an optical cavity — those perfect and near-perfect mirrors used to enclose it — a laser is created.

The kinds of lasers used for optical networking typically work by pulsing on and off. Not only does this allow a laser to produce a more intense burst of light energy when it emits, it also matches the binary nature of modern digital communications.

On and off matches the 1s and 0s used to represent digital data very nicely, in fact. Communication lasers use gain media that are carefully tuned to produce light in very tight frequencies ranges. How tight? Well, some optical lasers emit ranges as narrow as 20 nm (nanometers, one billionth of a meter). That is somewhere between 750 to 9,000 times narrower than the cross-section of an average human hair.

Fiber is the optical medium

Though some optical fibers are made of plastic, fibers used for optical networking cables that span any distances more than 10 meters or so are invariably made of glass. In fact, the glass that goes into fiber optic cable is very special: It's carefully fabricated to achieve total internal reflection (or as close as modern technology can get). This means that light beams always bounce inward whenever they encounter the outer layer of the glass fiber strand at the heart of the cable.

It turns out that fiber optic cables have all kinds of nifty technical advantages for communication. Signals travel across optical fibers with less loss over distance, which means that fiber cable segments can get ridiculously longer before inline amplifiers or regenerators must be inserted to restore signals to a pristine state. Whereas metallic network cable segments can seldom span much more than a mile (most types span much shorter distances), some types of fiber optic cable can span distances up to 60 kilometers (just over 40 miles) with a single segment. By itself, this helps to explain why fiber optic cable has become the networking medium of choice for long-haul communications.

But there are other reasons why fiber optic cable makes a good choice for high-volume, high-bandwidth communications. Optical fibers can carry huge amounts of information because they can accommodate an enormous number of light pulses of very short duration. A typical coaxial (copper) cable used for cable TV can deliver 1,000 or more TV channels and networking services with an aggregate bandwidth of less than 6.3 Gbps. A modern fiber optic cable can accommodate up to 80 channels at once, with data rates as high as 111 Gbps per channel.

That translates into 8.88 Tbps for a single cable, or more than 1,400 times the capacity of the coaxial cable used for cable TV. Ongoing experiments in the lab indicate that the technology is what limits fiber-optics-carrying capacity, however. Thus, most experts believe its maximum carrying capacity will keep increasing for the foreseeable future.

There's one more nice thing about fiber optic cable. It ferries light pulses, which are immune to electromagnetic interference (EMI, from electric motors, transformers, and other powerful electrical devices) as well as radio-frequency interference (RFI, from nearby signal transmitters of any kind, as well as many of the devices that also emit EMI). This makes fiber optic cable ideal when networking cables must traverse interference laden environments, such as factory floors, city streets, and so on.



There are two basic types of fiber, known as multi-mode fiber (MMF) and single-mode fiber (SMF). SMF usually comes in a yellow jacket or outer sheath, while MMF features an orange exterior. MMF is good for short distances of up to 500 meters and usually operates at a wavelength of 850 nm. SMF is good for long distances (10 kilometers and up) and uses wavelengths between 1310 nm and 1550 nm. Techniques such as coarse wavelength division multiplexing (CWDM) and dense wavelength division multiplexing (DWDM) allow multiple signals to occupy the same fiber cable. (These are described in detail in Chapter 3.)

MMF is cheaper than SMF but spans much shorter distances. SMF requires much more precision and entails more expensive equipment and set-up costs. But when light has a long way to travel, SMF is the only way to go!

Making fiber optic connections

If fiber optic technology has an Achilles' heel, it relates to connecting fiber optic cables to devices — or actually, mating up those cables to physical connectors that will attach to devices. The ends of a fiber optic cable must be cut precisely at the right angle (as perpendicular to the cable itself as

possible) and polished as smooth as modern technology will allow. This is the only way to ensure the most perfect fit possible with a connector that mates with the end of the cable and is used to attach it to some device.

The optical fiber used for SMF is significantly thinner than that used for MMF. Typical multimode fiber has a core diameter of 62.5 μm (micrometers), whereas single mode fiber is usually about 8–9 μm in diameter. This helps to explain why cable finishing and equipment costs for SMF cables are so much higher than those for MMF: there's significantly less room for mismatch or error of any kind!

Typical fiber optic connectors include the following six options:

- ✓ **LC:** These appear on the newest high density optical gear and are quickly becoming a standard connector type. The connector shown in Figure 2-2 is a duplex LC connector (with two strands of cable); simplex varieties are also available.




Figure 2-2: The LC is the high-density optical connector of choice.

- ✓ **SC:** These appear on lots of optical networking equipment and are standard for most older optical GbE gear. Check out Figure 2-3.




Figure 2-3: The SC is standard on older GbE equipment.

- ✓ **FC:** Includes a screw-on locking mechanism; commonly used with SONET and SDH equipment as well as optical test gear. See Figure 2-4.




Figure 2-4: The FC is common on SONET, SDH, and optical test equipment.

- ✓ **ST:** An older connector and never as popular as FC, but also used on SONET and SDH systems, as well as optical test gear. Shown in Figure 2-5.




Figure 2-5: The 1980's era ST was used on SONET, SDH, and test equipment.

- ✓ **E2000:** A mostly European connector not very common elsewhere in the world. Incorporates an automatic dust cover/safety mechanism that's proved of little practical use in the field. See Figure 2-6.




Figure 2-6: The E2000 is popular in Europe, seldom used elsewhere.

- ✓ **MT-RJ:** Initially introduced as the *de facto* high density optical connector until breakage issues appeared, and the LC connector took over. Some early high-density GbE switches use these. Shown in Figure 2-7.




Figure 2-7: The MT-RJ proved too fragile as a high-density optical connector.

Optical receivers/sensors

With a laser emitting light on one end of a fiber optic cable, a sensing or receiving device is required on the other end to detect incoming signals. Because so many forms of optical communication tend to be multiplexed — that is, to stack up multiple channels in different time slots or at different wavelengths — optical receivers must also act as demultiplexers as well as sensing incoming light pulses.

This puts the onus of precise wavelength filtering for incoming signals on the demultiplexer, because it sits at the front end of the incoming light pulses and decides which pulses belong to each channel. After that, the optical receiver may translate pulses of light into electronic digital data for local distribution or use in electronic form, or it may simply transmit channel contents optically to another optical device.

Fiber Optics 101

When putting together the pieces of a fiber optic network, numerous devices and items of equipment prove necessary. Here, we talk about how fiber optic communications work, then start digging into some of the other types of choices implementers of fiber optic networking are likely to face and make as they start building and testing.

First and foremost, communication is one-way only on any single fiber optic cable channel. There's a transmitter on one end of the channel and a receiver on the other end of the channel, and these two devices cannot change roles. Two separate fiber cables may be necessary to establish bidirectional communication in some cases. Otherwise, special equipment to set up separate channels on a single cable is

needed, where there's a transmitter on one and a matching receiver on the other end on one channel (from Point A to Point B), and vice versa on another channel (from Point B to Point A). The bottom line, however, is that on any single fiber optic communications channel communications are all one-way, and one-way only.

MMF versus SMF

Multi-mode fiber is coarser, cheaper, and spans shorter distances; single-mode fiber is finer, requires extreme precision and expensive equipment, and spans long distances. These differences virtually dictate that SMF is used for WAN applications, or for optical cable runs that exceed the 500 meter limit on MMF. This also means that SMF is the fiber optic cable of choice for communications carriers, telecom companies, and for heavy-duty high-capacity network links at metropolitan area or greater geographical scales.

A matter of loss

Some of the biggest differences between MMF and SMF cables also explain their respective best uses and niches. MMF uses larger, coarser optical fibers to transmit light from one end of the cable to the other. This gives light beams more room to bounce around and also causes more potential variation in the time it takes for beams to transit the cable from end to end (more room to bounce means more time for those signals that bounce a lot to make the transit, as opposed to those signals that bounce less).

This phenomenon introduces two types of distortion into signal transmissions known as *jitter* and *dispersion*, which is probably best understood as the amount of variation between the fastest transit time across the cable as compared to the slowest transit time. The greater that difference, the greater the distortion. Too much jitter can garble optical signals or make them unintelligible. MMF is subject to higher jitter rates, which explains why it's used only for relatively short distances.

There's also some materials science involved in calculating the speed of light through any medium. Light travels fastest through a vacuum because there's nothing to impede its travel. Even over glass fibers, the speed of light drops, as photons must push their way through glass molecules to move from one end of the cable to the other. MMF cables typically use coarser materials as well as larger fibers, and the increased tolerance for impurities also means that light moves more slowly through MMF than through SMF optical fibers — not just because of higher external reflectivity and shorter bounce paths, but also because of the higher number of nontransparent molecules that photons must find their way around.

Optical receivers

Optical receivers are best judged in terms of their overall sensitivity, which really translates into how low light levels can get before they become undetectable or indecipherable.

Optical receivers come in two basic types:

- ✓ **PIN:** PIN is an acronym for the construction materials in a certain type of optical receiver called a *photodiode*. P stands for P-type material, which is a type of silicon that is positively charged, I stands for integral silicon which is neutral in charge, and N stands for N-type material which is negatively charged. When light strikes PIN materials, it causes electrons to migrate to the negative silicon, and positive charges to increase on the negative silicon, thereby inducing a current increase. Photodiodes can be tuned to be sensitive to specific wavelengths, if needed.
- ✓ **APD:** An *Avalanche Photodiode* is a PIN receiver with an extra semiconductor layer across which high voltage is applied. When light hits the APD, an “avalanche” of impact ionization occurs in that layer, allowing the APD to deliver a more powerful signal. This means an APD is more sensitive than a PIN receiver, but it also means that APDs cost more because they use higher-voltage power supplies and produce more electronic noise.

More expensive, longer range optical networking gear is more likely to use APD receivers, whereas less expensive, shorter range equipment is more likely to use PIN receivers.

All optical networking?

When optical signals arrive at the end of an optical fiber, a receiver picks them up. The receiver will generally pass on the information included in those signals, usually on a per-channel basis, either for retransmission or for conversion into electronic format for use on non-optical network segments.

When signals will stay on an optical path, there are two options for forwarding them from the receiver that accepts them as input to the laser that produces them as output. They may be forwarded in electronic form, which requires converting the signal from optical to electronic form for transmission from the receiver to the laser (known as OEO for *optical to electronic to optical*), or they may be forwarded in optical form, which keeps the signals all optical all the time.

Lasers and receivers cost money, so it's never a good idea to use more of them than you really need. Some all-optical concepts and equipment actually work, but they tend to be quite expensive. This also explains why you still can't buy an all-optical computer on the open market, even though engineers have been talking about such devices for the past two decades, and have even built numerous prototypes

Best practice dictates that deployment use amplifiers (OEO) rather than repeaters and small optical switches (all-optical) when moving optical data from one fiber to the next. Not only does this cost significantly less, but repeaters also regenerate signals and restore them to their original fidelity as they convert optical signal to electronic equivalents and then convert those electronic signals back again into optical ones.

One vexing issue with all-optical transmission occurs when bouncing light off a mirror (a common technique for turning light output from one source into an input for another source). Some loss of signal strength, coherence (the opposite of jitter), and fidelity is bound to occur as this happens.

Amplifiers and repeaters

Fiber optic cables may span long distances, but even they are subject to *attenuation* (a weakening of signals carried as the distance increases). Optical amplifiers take weaker incoming

light signals on the input side and use special technology to increase the power of light signals emitted from the output side. When fiber optic cables must span really long distances (60–100 kilometers per segment is a typical limit), amplifiers must be inserted inline to keep signals sufficiently strong to be understood when they arrive on the receiving end of the fiber link. This is one way to address the loss of signal strength as signals move across a fiber optic medium.

Another way is to use a device called a repeater between fiber optic cable segments. A *repeater* is an OEO device and not only strengthens the input signal, but also recreates the original binary data stream as closely as possible. Proper spacing of repeaters enables extremely high fidelity from end to end. Amplifiers, on the other hand, do nothing to improve signal accuracy or fidelity, they only boost signal strength.

Chromatic dispersion

Chromatic dispersion occurs in light signals much like the way a prism separates colors in light into separate, distinct wavelengths. Chromatic dispersion becomes a problem as networking speeds increase. At 10 Gbps and higher, it can be serious. The effects of chromatic dispersion increase at the square of the speed so that 10 Gbps is 16 times as prone to this problem as is 2.5 Gbps, and 40 Gbps is 256 times more prone than 2.5 Gbps.

Compensating for chromatic dispersion requires using dispersion-compensating fiber that shifts all the component wavelengths back together. Essentially, this material resynchronizes the wavelengths in the light it carries. The transmitter is the primary factor in determining chromatic dispersion. This explains why tight, narrow wavelengths are used for high-speed optical communications: the tighter the wavelength, the fewer the colors it can carry and the less dispersion becomes possible.

Polarization mode dispersion

At extremely high speeds, only single mode fiber will do. But, single mode fiber is designed to transport light wavelengths

in the form of two perpendicular waveforms, also known as *polarization modes*. Looking at such modes in a fiber, you would see two perpendicularly crossed lines like the letter X.

Polarization mode dispersion (PMD) occurs because the glass in optical fiber is never 100 percent pure silicon. Minor variation in fiber content and density can disturb the relationship with polarization modes so that one travels at a slightly different speed from the other. At speeds up to 40 Gbps, this isn't a problem, but at 40 Gbps and higher, PMD causes the same kinds of signal losses and distortions that occur at lower speeds with chromatic dispersion.

At present, significant research is underway to learn how to compensate for PMD. So far, all solutions developed are very expensive. Because other technologies are cheaper (especially wave division multiplexing and 10 Gbps optical Ethernet interfaces, both discussed in Chapter 3), no marketable solutions are expected any time soon.

Chapter 3

Building Bigger Networks

.....

In This Chapter

- ▶ Multiplying colors and bandwidth
 - ▶ Introducing wavelength division multiplexing
 - ▶ Working with coarse wavelength division multiplexing
 - ▶ Using dense Wavelength division multiplexing
 - ▶ Distinguishing between coarse and dense versions
-

Today, nearly all telecommunication spends a significant portion of its lifespan in the form of fast-traveling photons. How exactly does this work? Fiber optics, of course! But, that's an incomplete answer that doesn't much address the "how" and the "exactly" parts.

To tackle those details, we frame the answer in an ongoing example that illuminates the principles of light-based communication. In this chapter, we shed some light on the subject of communicating with light signals.

Double, Triple, and Tenfold Bandwidth

Imagine two distant towns surrounded on all sides by mountain peaks where neighboring townsfolk can see each other from atop an elevated perch. Telephony hasn't quite taken hold in either town, but the local townspeople have discovered that they can communicate with each other with beams of bright light.

When Alysia wants to contact Braedon, she sends a series of long and short flashes of light. Okay, we use this simplistic example to show how optical equipment communicates, but instead of using fresh mountain air as a carrier, it uses optical fiber. Alysia and Braedon decide to use this floodlight method as an effective primary means of communication instead of relying upon letter writing or carrier pigeons.

Along comes Alysia's neighbor Ceci, who recognizes this strategy as a practical means for communicating with Damon, who also lives near Braedon. Ceci also decides to buy a floodlight and stand atop the same mountain peak as Alysia and communicate more or less at the same time Alysia and Braedon are communicating.

Cross-talk communication

If ever you've listened in on a conversation between two people where both respond at the same time, one talking over the other, you know how difficult it can be to make any sense of anything. Communicating like this creates contention because there's no reasonable way for Damon to distinguish which flashes are meant for him and which are meant for Braedon.

Both Alysia and Ceci appear to communicate from the same point because they're at the same position and using the same kind of floodlight. Messages get garbled and nobody understands anybody. Alysia misinterprets Damon's invitation to dinner as Braedon's decision to break up, and Braedon misconstrues Ceci's signaling as a marriage proposal. What a mess!

This is similar to what happens when two optical interfaces use the same fiber connection. Unlike the CAT5 cable-based equipment you're probably most familiar with, fiber optics communication doesn't support mechanisms to detect when another machine is transmitting, nor can it back off and retry when conversations collide.

Eliminating cross-talk

If there were another mountain peak nearby, Ceci could use that instead. This is like having a second optical fiber to

interface and interact with another receiver. Perhaps the alternate peak is inaccessible and has no clear line of sight to Damon, so Ceci insists on using the same location as Alysia — at the same time.

During telecom's booming peak period, lots of optical fiber got deployed nationwide making it easier to use individual fibers for individual subscriber services. Smart providers lease fibers to other companies at a huge premium to cover the costs of operating a fiber optical network. These people also understand that optical fiber is limited and may someday become scarce.

As Alysia, Braedon, Ceci, and Damon improve their signaling skills, they become much faster at flashing lights and agree to take turns transmitting — especially after that recent embarrassing incident.

By flashing light more quickly, all parties can send the same amount of information in a shorter time span and share the same hours for remote communication. This is somewhat like Time Division Multiplexing (TDM), which involves sending smaller chunks of data at higher rates of speed, interspersed with other chunks of data sent likewise, to achieve greater overall throughput.

Shortcomings of TDM

The same problem that Alysia and Ceci experienced recurs, however, when newfound lovebirds Elsie and Francis join the fray. Oh, boy — here comes trouble all over again. Alysia and Ceci could try flashing light three times faster than before, which is like upgrading to higher-speed fiber optics — all at a much higher cost premium, of course — but that's too much work.

Imagine that Alysia and Braedon can take advanced signaling courses at an added cost leaving Ceci and Elsie to utilize beginner-level signaling. That's analogous to using fiber optic media, such as the synchronous optical network (SONET) optical carrier level 3 (OC-3), synchronous digital hierarchy (SDH), synchronous transport model 1 (STM1), or Gigabit

Ethernet (GbE). Advanced communications courses, the kind Alysia and Braedon are interested in, are equivalent to 10 GbE (10G) to 40G fiber optics, which are very expensive. We cover these technologies later in this chapter.

Eventually, Alysia and Braedon will hit an upper limit on how fast they can flash a floodlight to keep chatting from their mountaintop positions. Ceci and Damon aren't far behind them in ability and are now struggling with Elsie's and Francis' flirtatious exchanges. There must be a better way for all parties involved. Something similar happens when 10G optics encounter chromatic dispersion or when 40G optics encounter polarization mode dispersion (PMD).



Chromatic dispersion is the broadening of light pulses caused by differing propagation speeds for various light wavelengths. PMD is an inherent property of all optical media caused by the differences in propagation velocities for light in an optical carrier network.

Another Way of Multiplexing

In a flash of brilliance, Braedon gets a bright idea. If he and Alysia use a colored lens, they can easily differentiate themselves from the other couples regardless of how fast or slow they transmit data, and irrespective of the method everyone uses for signaling. Alysia and Braedon instantly put this method to use with immediate success.

Alysia and Braedon now communicate using carefree flashes of red, followed shortly by Ceci and Damon in green, and Elsie and Francis using a light shade of blue. It's this crucial distinction that makes Wavelength Division Multiplexing (WDM) work — by using different colors of light. Each party “tunes” its particular lights to some distinctive color to avoid confusion among recipients.



WDM is a method for multiplexing multiple optical carrier signals on a single fiber using laser that emits separate and distinct wavelengths (colors) to carry different signals.

Creating optical filters

In terms of WDM, Alysia and Braedon have created optical filters. This enables optical communications devices to separate wavelengths of light using a shared medium. In WDM, there are filters on both sides. One side will add all wavelengths together on a fiber — a process known as multiplexing — and the other side decouples and separates wavelengths — a process called *demultiplexing*.

When Alysia and Braedon encountered problems communicating at the same time as Ceci and Damon, they discovered TDM as a cheap and effective way to share the same medium (in their case, fresh mountain air).

Then came Elsie and Francis, inspiring Braedon to think a little brighter. That's how he discovered WDM, which is a little more costly to implement with the addition of colored lens filters. If they could all just use the telephone system or high-speed networking technologies, none of this would be necessary — but then, we'd have no way to illustrate how optical communications work.

Types of WDM

In this section, we talk about the two types of WDM currently in existence today: coarse WDM (CWDM) and dense WDM (DWDM). Backwards as it may seem, dense WDM came well before CWDM, which appeared only after a booming telecommunications market drove prices to affordable lows.

What Braedon discovered with lens filter coloration is analogous to CWDM. Had Alysia and Elsie opted for differing shades of red instead of entirely different colors, they would have demonstrated DWDM. Whereas CWDM breaks the spectrum into big chunks, DWDM dices it finely, as shown in Figure 3-1.

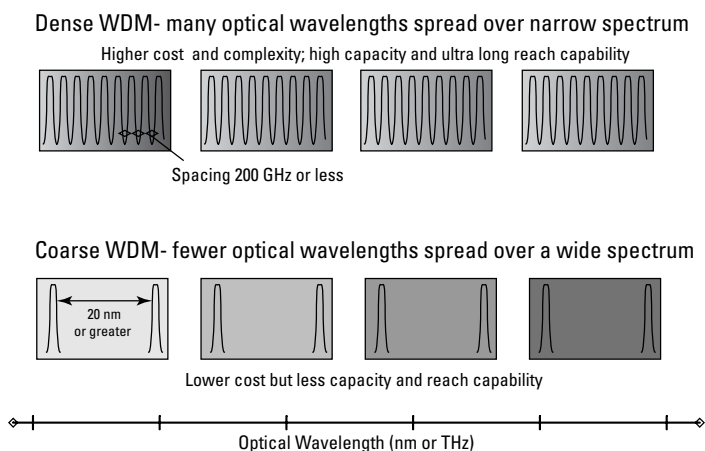


Figure 3-1: Comparing CWDM and DWDM frequency bands.

DWDM fits 40-plus channels into the same frequency range used for two CWDM channels. Can you name 40 different colors? Don't strain yourself — even a professional artist would take a while. The DWDM C-band is about 40 nanometers (nm) wide, which means that you must not only name 40 different colors but they must all be shades of red, as well. How tiring!



For our purposes, a *C-band* or color band represents a certain range of light wavelengths used for communications. Its variations refer to various broadcast ranges for different communications technologies; you need only understand it as way of describing color spectrums.

Naming WDM channels

When optical engineers discuss WDM, they use abbreviated formats to refer to individual channels. Instead of “Channel 1551,” they say “Channel 55.” This isn't done out of laziness or bad math — well, an engineer might be lazy, but that's irrelevant — but for a perfectly valid reason. Identifying a laser's output at 1551 nm references a particular wavelength.

Earlier in this chapter, we told you that CWDM is defined by wavelengths. These wavelengths are spread 20nm apart (such as 1451 nm, 1471 nm, and 1491 nm), called *channel spacing*. Each wavelength is represented by some odd number written as 1xx1, where each x represents a single digit. Since each wavelength is surrounded by enclosing 1s, these can simply be omitted using a short-form Chxx notation. In the future, forthcoming versions of CWDM may have channel spacings of 10 nm or 5 nm and be annotated accordingly.

Channel spacing frequencies

DWDM is defined in terms of frequencies. There are several types of DWDM with channel spacings that are 50 GHz, 100 GHz, and 200 GHz apart. Smaller spacings fit more channels onto a single fiber, but cost more to implement and operate.

Equipment with 200 GHz channel spacings uses frequencies at 192.3 THz and 192.5 THz. Here again, we count in odd numbers. Because all channels use a 19x.x format, we can drop the static bits of information (in this case, 19) and identify these channels as Ch23 or Ch25.



CWDM and DWDM are different technologies with entirely different properties. Thus, CWDM channel 1571 nm (Ch57) is not the same as DWDM channel 195.7 THz (Ch57).

Distinctive CWDM differences

CWDM can — in principle — match the basic capabilities of DWDM but at lower capacity and lower cost. CWDM enables carriers to respond flexibly to diverse customer needs in metropolitan regions where fiber may be at a premium. However, it's not really in competition with DWDM as both fulfill distinct roles that largely depend upon carrier-specific circumstances and requirements anyway.

The point and purpose of CWDM is short-range communications. It uses wide-range frequencies and spreads wavelengths far apart from each other. Standardized channel spacing permits room for wavelength drift as lasers heat up and cool down during operation. By design, CWDM equipment is compact and cost-effective as compared to DWDM designs.

Distinctive DWDM differences

DWDM is designed for long-haul transmission where wavelengths are packed tightly together. Vendors have found various techniques for cramming 32, 64, or 128 wavelengths into a fiber. When boosted by Erbium Doped-Fiber Amplifiers (EDFAs) — a sort of performance enhancer for high-speed communications, these systems can work over thousands of kilometers. It's something like steroids for high-intensity athletes, only much lighter (and legal)!

Densely-packed channels aren't without their limitations. First, high-precision filters are required to peel away one specific wavelength without interfering with neighboring wavelengths. Those don't come cheap.

Second, precision lasers must keep channels exactly on target. That nearly always means such lasers must operate at a constant temperature. High-precision, high-stability lasers are expensive, as are related cooling systems and costs.

CWDM and DWDM scenarios

CWDM doesn't span long distances because its light signal isn't amplified, which keeps costs down but also limits maximum propagation distances. Vendors may cite working ranges of 50 to 80 kilometers, with 160 kilometer distances achievable through signal amplifiers. CWDM supports fewer channels and that may be adequate for metro carriers who prefer to start small and expand later as demand increases.

Here are four typical scenarios for WDM deployment:

- ✔ **Case 1:** Single channel systems, one channel per fiber. Each fiber carries one wavelength.
- ✔ **Case 2:** 8 channels across a single fiber pair.



✓ **Case 3:** 16 channels across a single fiber pair.


✓ **Case 4:** CWDM using 16 or 32 channels across a single fiber pair.

Wavelengths are typically spaced at 200 GHz intervals for metropolitan applications.

The four scenarios described in the preceding list are depicted in Figure 3-2.


Non-amplified signaling systems keep entry costs down and can still retain high loss tolerance. Whenever a non-amplified signal is used, there is a trade-off between capacity and distance. Either you make long networks with fewer nodes or smaller networks with lots of nodes.

Case 1: G.652 or G.652.C fiber pairs with single channel systems




(1 λ per fiber)

Case 2: G.652 or G.652.C fiber pairs with 8-channel CWDM system




(8 λ per fiber)

Case 3: G.652 fiber pairs with 16-channel CWDM system



(16 λ per fiber)

Case 4: G.652 or G.652.C fiber pairs with DWDM system



(16 to 32 λ per fiber)

Figure 3-2: Diagrams for four typical CWDM deployment scenarios.

Core Terminologies and Principles

Okay, it's time to get down to brass tacks. When you start interfacing with equipment vendors and service providers, there are some specific terminologies and principles you should know.

In the olden days, all phone conversations were simply analog signals traveling over copper wire. This worked fine when telephone lines connected city to city but had several unique problems. As an analog signal travels greater distances, it suffers *attenuation* — a fancy way of saying it grows weaker. Imagine for a second that Alysia and Braedon instead shouted from each mountain peak during the day and flashed beams of light by night. During the daytime, their conversations required shouting because of low volume. The “signal” could have been amplified, perhaps through the use of bullhorns, but some forms of amplification introduce line noise in the form of static or hissing sounds.

When analog signals are transported closely together, they tend to have a mutually interruptive affect called *crosstalk*. You've probably been on an older cordless phone when suddenly you pick up a neighbor's conversation — that's crosstalk. It also means that at the far end, your voice comes back to you on the return path causing an unappealing echo.

Analog signal analysis

Now, we visualize how an analog signal travels. We start our conversation with an analog signal that becomes divided into time slots and measured at the start of each slot. Each time slot gives a number that's transported through the digital network. On the receiving end, we connect time-slot measurements together at the very beginning of each.

Signal quality can vary based on the size of the sampling — the width between rising and falling curves of a sample — and the number of heights each bar has. A single phone call has 7 or 8 bits (meaning 128 or 256 levels) and is sampled 8,000 times per second. A high-end CD player with 24-bit sampling

has 2^{24} or 16.8 million levels, and samples 44,100 times per second. An analysis of voice signals appears in Figure 3-3, and shows how narrow the audible sound spectrum really is.

Digital networking transmits data in numeric form. As long as you can read a number, the number can be transmitted and completely restored without introducing noise. This is a process known as *repeating* or *regenerating a digital signal*. Digital signaling like this isn't prone to actual crosstalk, as when Alysia and Ceci began shouting or flashing light beams, because every processed digital signal filters out line noise.

More important, we can deliver numbers faster than they can be read. If we can read two numbers once per second but transmit them in half a second, then we can use the same line to transmit both numbers. This is the basis for TDM, introduced at the outset of this chapter.




Figure 3-3: Dissecting an analog voice waveform.

Key TDM provisions

AT&T developed synchronous optical networking (SONET). This format originated at about 50 Mbps and quickly grew to multiplexing three channels into what's called an Optical Carrier level 3 or OC-3.

The synchronous digital hierarchy (SDH) is a separate standard based on *circuit mode communication*, meaning that each connection achieves a constant bit rate (CBR) and constant delay. SDH may be utilized simultaneously among several different ISPs on the same optical fiber without ill effect to each other's traffic load and without borrowing free capacity.

Both SONET and SDH are pure TDM protocols offering permanent connections and packetless physical layer transmission. Although SONET is widely used throughout North America,

SDH is used throughout much of the rest of the world. SONET and SDH are nearly identical optical methods that use a common framing scheme where the major difference between them is the numbering used to identify interface speeds.

Ethernet at fiber speeds

Within the past two decades, Ethernet technology has reached a level of operation that rivals SONET. The 10 Gigabit Ethernet (10G) standard is becoming commonplace, allowing users to transport Ethernet signals at a rate that's comparable to existing SONET connections. 40G and even 100G Ethernet standards are in the oven and will eventually become common.

Like SONET and SDH, Ethernet is also a frame-based (that is, packetless) transmission protocol that is about as prevalent as the packet-oriented Internet Protocol (IP). Unlike IP, Ethernet has a well-defined physical layer specification that maps directly onto fiber.



For more information on how Ethernet maps into the fiber standard, see IEEE specification 802.3an (a free PDF version is available online at <http://standards.ieee.org/getieee802.3/802.3.html>).

Grading Ethernet types

Ethernet can be made carrier-grade, but that requires additions and enhancements. The five attributes discussed here bring LAN-based Ethernet to standardized service levels capable of reaching carrier-grade quality.

- ✓ **Quality of service:** Carrier Ethernet offers a range of granular bandwidth options. Defined QoS attributes enable assured service-level agreements (SLAs) to protect against frame loss, network latency, and delay-variation minimums.
- ✓ **Reliability factor:** Traditional LAN-based Ethernet is a best-effort connectivity mechanism. Carrier Ethernet rapidly detects and recovers node, link, and service failures in less than 50 milliseconds.

- ✓ **Scalability factor:** Ethernet is already available in 10 Mbps to 10 Gbps varieties. A single interface can scale from 10 Mbps to 1 Gbps via software control; Quality of Service (QoS) guarantees seamless delivery over a common infrastructure and spans broad geographies through different providers.
- ✓ **Service management:** The most critical attribute is the ability to diagnose, monitor, and centrally manage the network using standards-based, vendor-neutral tools to rapidly provision services, troubleshoot connectivity issues, resolve faults, and measure performance characteristics.
- ✓ **Standardized services:** There are two defined standard service types for the delivery of Ethernet:
 - **E-Line:** A point-to-point Ethernet virtual connection (EVC) between two user networks works through Ethernet private line (EPL) or Ethernet virtual private line (EVPL) services.
 - **E-LAN:** A multipoint-to-multipoint Ethernet virtual connection between user network interfaces on a Carrier Ethernet network.

Without these additions and enhancements to the Ethernet standard, there would be no translation into an SLA-bound, carrier-grade format. Carrier-grade Ethernet extends Ethernet into the next generation of super high-speed metropolitan networks. Scalable secure segregation, fault tolerance, and multi-service capabilities correct deficiencies with standard Ethernet that would otherwise restrict its transition into metropolitan area networking.

Doing More with the Same Fiber

Because of inspired advances in optical communications, Alysia and Braedon can communicate at very high rates of speed. Since the development of coarse and dense wavelength gradation, they no longer suffer from crosstalk with Elsie and Damon — or any other townsfolk — and can cram more conversations into the same channels.

In Chapter 4, we look into making the most of your high-speed optical network including traffic grooming, network redundancy, and some useful protection schemes.

Chapter 4

Making the Most of Your Optical Network

In This Chapter

- ▶ Understanding traffic grooming
 - ▶ Digging into traffic grooming strategies
 - ▶ Building network reliability and redundancy
 - ▶ Exploiting ring and mesh network topologies
-

Two primary methods for handling significant temporal overlaps in communications are time division multiplexing (TDM) and wavelength division multiplexing (WDM), each of which incorporates various derivatives. In this chapter, we look at how optical networks manage to cram even more communications potential into the same optical medium to better utilize high-speed fiber optic capacity.

In the previous chapter, we illustrate the basic operation of optical networking using a simplified analogy with colored lights and colored lenses. Here, we talk about complications that can arise when multiple sources of light (in our example case, bright floodlights) transmit signals at the same time.

Grooming Is More than Good Hygiene

WDM is used to expand capacity in optical networking (discussed in detail in Chapter 3). In a WDM network, each

optical fiber communications link can carry high-rate traffic at varying wavelengths. Multiple channels can be provisioned through a single fiber link.

When we talk about traffic grooming, it's not what you might initially think. There is no effervescent scrubbing or fine-bristle brushing involved. In terms of networking, traffic grooming defines a process for utilizing network bandwidth more effectively and efficiently.

Increasing optical carrier capacity

Drawing on our example from Chapter 3, recall that our mountaintop communications hit certain limitations as more people and more sources of light became involved at the same time. With only a single communications medium available — in that example, open air — there are only so many people and so many conversations that can occur at any one time. However, clever thinking and careful implementation enabled more people to communicate using different colors of light so that everybody could tune in to their intended messages and ignore other communications concurrently underway.

As more townspeople become involved in the mountaintop light show — particularly as each individual communicates with more than one person — the limitations and inefficiencies of TDM and WDM become more apparent. Now, imagine that instead of relying on several separate sources of light, everyone learns to communicate through a single, much larger light source.

Cleaning and preening

Traffic grooming is a technique that groups multiple smaller telecommunications sub-streams (such as each individual conversation between Alysia and Braedon, Ceci and Damon, and so forth) into larger cohesive units that can be processed in a single information transfer.

For networks that use TDM and WDM, two traffic flows destined for a common location can both use the same wavelength (or channel) to communicate. This enables both flows to pass through a single optical add-drop multiplexer (ADM), which acts like an on- and off-ramp for high-speed networking.



An ADM combines (multiplexes) several lower-bandwidth streams of data into a single higher-bandwidth optical bit stream (which enables it to accommodate multiple conversations or communication sessions at the same time). As an acronym, ADM should not be understood as an alternative to TDM or WDM but as an integral aspect of any modern optical fiber network.

You may also hear an ADM referred to as line-terminating equipment (LTE). It is also frequently the costliest component in any WDM's construction. This opens our discussion to various traffic grooming techniques that ADMs can use.

Aggregate communications channels

Laying optical fiber and establishing optical nodes is an expensive proposition. By itself, this is a darn good reason why engineers are always interested in making the most of what bandwidth is already available. Traffic grooming seeks to extract flows within a given stream and place them into other flows destined for a common location.

Imagine for a moment that Alysia acquires another friend named Glenda in the same town as her friend Braedon. Meanwhile, Ceci develops a budding friendship with Harry in another town. At the same time, other mountaintop communicators establish “flows” of communication to other townspeople in towns where Braedon and Glenda reside.

Instead of creating a confusing collection of light beams emanating in every direction, the townsfolk devise a method for beaming a single light that combines all communications for a single town into a single source. All communications

destined for Braedon's town are aggregated and sent together, while all conversations meant for Glenda's town are combined and sent using a separate and unique light source. This is how traffic grooming works, at a very superficial level. There's a lot of fancy math and communications algorithms involved (which helps explain why the necessary gear is so expensive), but traffic grooming is worth using because it enables fiber optic links to be used more efficiently and effectively.

Within an optical network, traffic flows streaming from Los Angeles destined for Washington, D.C. are extracted and placed into a single large traffic flow. This flow is directed onto a port connected to fiber optic cable with a path to D.C., thereby enabling maximum utilization of that port and minimizing the number of ports and fibers required to transport traffic between L.A. and D.C.

Traffic Grooming Techniques

WDM is widely used to expand optical network capacity. Each high-rate fiber link carries traffic at many different wavelengths where multiple channels are created within a single fiber.

Two basic architectures are employed for WDM networking: *ring* and *mesh*. Most optical networks in use today are built on ring topologies, but carriers are increasingly moving to mesh architectures for next generation networks. Mesh networks have a compelling cost advantage over ring networks, are more resilient to network failure, and more accommodating when changes in traffic or demand levels occur.



We discuss ring and mesh topologies in more detail later in this chapter.

Traffic grooming, routing, and wavelength assignment are among the more important issues to consider when designing an optical mesh network. The problem with traffic grooming and routing within mesh networks lies in determining how to most efficiently route traffic from source to destination, while also combining lower-rate or *sub-wavelength* traffic into single wavelengths for high-bandwidth transmissions.

Electrical grooming techniques

A key component that enables traffic grooming in mesh networks is an Electrical Cross-Connect (EXC) coupled with an Optical Cross-Connect (OXC). This hybrid architecture handles data switching and traffic grooming at the electrical level while simultaneously multiplexing, demultiplexing, and switching optical traffic streams at the optical level.

An OXC is a device used by telecommunications carriers to switch high-speed optical signals in a fiber optic network. There is more than one way to achieve this: electronically and photonically. With electronic cross-connect conversion, an optical signal is translated from optical signals to electronic signals for grooming and back to optical again. Such OEO designs (discussed in Chapter 2) are subject to one important limitation: electronic circuits limit the maximum bandwidth for each signal.

Envision an intermediary device capable of intercepting, interpreting, and interconnecting light signals — both electronically and optically — between mountain peaks. Alysia's light beam transmission to Braedon is intercepted by this electronic device that converts optical signals to electronic signals. These electronic signals are then switched by an electronic switching module, then converted back into optical signals before being reissued to the fiber network as optical light again. By design, the OEO conversion from optical to electronic signaling enables signal quality monitoring and has the added advantage of regenerating signals, which keeps nodes free of attenuation and dispersion. Electronic OXC is also known as *opaque OXC* because it doesn't use light from end to end.

Look in the Mirror, Mr. Wavelength

Now, imagine that this cross-connect signal switching occurs within an entirely optical intermediary device. There is no electronic conversion of Alysia's optical signals to Braedon — everything remains entirely light-based. This type of OXC switch is called *transparent OXC* or photonic cross-connect

(PXC). An optical signal is demultiplexed and the resulting wavelengths are then switched by optical modules. It's all done with mirrors (usually).

As this conversion completes, signals are multiplexed onto output fibers by optical multiplexer devices and thereby retain their data rate and protocol transparency.

Reconfigurable optical ADMs

Earlier in the chapter, we talk about ADMs enabling traffic aggregation for metro and regional networking. Now, we turn to reconfigurable optical ADMs or ROADMs. ROADMs are designed to deliver greater flexibility to WDM networks — particularly DWDM networks — by enabling dynamic, transparent optical wavelength add-and-drop functionality.

ROADMs add considerable agility and robustness to optical network architectures, vastly improving service velocity and lowering ownership costs. They automate and simplify the planning and configuration of optical networks by allowing any of the 40 wavelengths on a fiber to be added, dropped, or expressed optically in any combination through a node.

This technology enables traffic to pass through network locations transparently without any need for conversion from optical signals to electrical signals and back to optical again.

Is This Thing Reliable?

With so much light and light-speed transmission, a valid question arises: Is this thing reliable? Humans are notoriously fallible, so it only seems logical that anything made by humans must also be inherently flawed.

As neighboring towns around Alysia and Braedon begin to pick up on their innovative light-based communications, more townspeople become involved. Alysia and Braedon are onto something big — and they begin communicating with other people in towns that aren't directly in their line of sight. As the system becomes more complex, problems that must be addressed begin to crop up.

Network redundancy

The goal of redundant network topology is to eliminate downtime associated with or caused by any single point of failure. All networks need redundancy to enhance reliability, especially high-speed carrier-grade networks (Figure 4-1 shows a network diagram with duplicate ADMs and fiber for complete redundancy).

Network reliability is achievable only through reliable network equipment designed with fault and failure tolerances. High-speed carrier-level networks are designed to heal themselves so that any faults and failures get bypassed quickly and automatically. In a redundant configuration, devices use numerous interconnections between network nodes to provide fail-over connections between endpoints.

When an ADM fails or a fiber is cut on a redundantly connected network, all connectivity is preserved by routing traffic through an alternate path that can deliver traffic to the same endpoints. Each node has two or more pathways at its disposal to maintain connectivity.




Figure 4-1: With the right optical components in place, a network can heal itself following a break in any single segment, often without service interruption.

Protection schemes

Traffic protection takes the form of rerouting or diverting traffic to other pathways when multiple failures occur. This critical functionality enables delivery of real-time services — streaming voice and video applications, for example — and provides the robustness necessary to support all kinds of demanding applications.

Applicable sources include grid computing, business connectivity, and financial transactions. In fact, any situation where various forms of outage could incur heavy financial losses on an hourly or per-minute basis is probably a situation where companies and organizations will readily understand the need for and value of such protection schemes.

Optical Network Topologies

A network's topology expresses the arrangement of and mappings between its physical and logical interconnections. The common network topology you may be familiar with is the local area network (LAN) topology that exhibits both physical and logical aspects.



Physical topology refers to the physical elements of network design that include devices, locations, and cable installations. A logical topology is defined by transitory paths that data takes through a physical network, which may differ from the physical design itself (an excellent case in point is IBM's token ring, which is actually a star-based topology that uses ring-based traffic routing and delivery mechanisms).

Any given node on a network, be it a LAN or a WAN, has one or several links to other nodes on the network. The mapping of these nodes and links produces a geometrical arrangement that describes the network's physical topology. Likewise, mapping the flow of data between endpoints describes the network's logical topology. There are many other network topologies in use today including bus, star, tree, point-to-point, hub-and-spoke, and hybrid (a mix of the preceding types). The topologies most relevant to optical networking are ring and mesh.

Assume for a moment that Alysia's and Braedon's placement atop each mountain matches their physical locations. This simply means that where Alysia and Braedon are actually situated defines their physical topology. How they manage to communicate defines their logical topology. Increasing involvement among more parties and their positional arrangements redefines the network topology they create — both physically and logically.

Okay, now we take a look at optical ring and mesh topologies.

Ring topology

In a ring network topology, each node is connected to exactly two other nodes, forming a circular pathway for network signals, where the final node in a network also connects to the first node in the same network. Traffic travels around this ring and each node handles every packet, as shown in Figure 4-2.

Because there's a one-way path between any two nodes, ring networks are easily disrupted or disturbed by failure or problems with any single link. Node failures and cable breaks can potentially isolate every node attached to the ring.




Figure 4-2: The ring network topology illustrated.

Ring topology exemplified

Imagine that Alysia sends a message to Braedon that gets relayed to a mutual friend of theirs named Ian. Ian lives in a town that is visible to Braedon but is shielded from Alysia's view, so she's forced to transmit messages through Braedon using the same light-based communications they've come to depend upon.

What happens if Braedon misses a message or abstains from the three-party link they've established? Ian then no longer receives Alysia's messages, and Braedon's got some explaining to do. Alysia isn't happy, and Ian is left wondering why nobody is telling him anything. What a predicament!

Ring topologies make for orderly network layout among every connected device. Rings exhibit great performance under load and include no requirement for central management of connections between end-point stations. However, a ring topology will be compromised by any end-point failures or bad network ports. Also, any change, addition, or relocation of end-point devices affects the entire network.

Luckily, there are ring protection methods such as line and path switched rings. These topologies evolved from telephony networks as an essential part of the SONET standard. Nodes on a switched ring monitor the health of optical transmission and in the event of a failure or cable break, switch a loop-back around the problem, limiting the damage. But switched rings lose steam after one failure.

Mesh topology

In a *mesh network topology*, continuous connectivity and reconfiguration around broken or blocked paths occurs by hopping from node to node until the final destination is reached. Mesh networking doesn't follow the same chain of inter-node dependency as a ring topology and avoids problems related to ring arrangements. The mesh shown in Figure 4-3 includes two to four links per node.




Figure 4-3: The mesh network topology illustrated.

In fact, mesh networks differ from others in that all requisite parts can interconnect to one another through multiple hops, and they are generally immobile. Mesh networks are also self-healing in that the network continues to operate even when multiple nodes break down or connections fail. As a result, a reliable network formation is arranged: As long as some path is available between sender and receiver, traffic gets through.

Mesh topology exemplified

Let's revisit the arrangement between Alysia, Braedon, and Ian from our discussion of ring networking. This time, a fourth party (Jocelyn) gets involved. Jocelyn has a clear line of sight between Alysia and Ian and is positioned at approximately the same distance between them as Braedon, but she's on a different mountain peak.

If and when Braedon becomes unavailable, Alysia can rely on Jocelyn to relay messages to Ian. Likewise, Braedon can take over whenever Jocelyn is unavailable. Alysia can now use multiple pathways — she can communicate through either Braedon or Jocelyn as she sees fit.

Chapter 5

Ten (Or More) Views of the Future

.....

In This Chapter

- ▶ Taking an overview of telecomm
 - ▶ Examining circuit-switching networks
 - ▶ Peering into packet-switching networks
 - ▶ Converging on a unified view of the network
 - ▶ Gazing toward tomorrow
-

Today's high-speed communications networks would be far different were it not for the rich story behind telephony. Mankind began communicating at great distances first by using smoke signals, then telegraph, postal services, and eventually over radio signals. Although telephone literally means “sound at a distance” or “sound from afar,” humans have always sought better ways to keep in touch.

Various simple and complex codes and symbols have been developed to further enhance long-distance transmissions. Carrier pigeons retired with the advent of long distance communication. But, the telecommunications era begins with the transmission of analog voice signals.

In this chapter — you guessed it — we take a look at yesterday, now, and what's yet to come.

Telco and Data Networks: Strange Bedfellows

Moving information over wires actually precedes the invention of the telephone. The telegraph and the fax machine beat the phone to the punch in the nineteenth century. Can we argue that networking precedes telephony? Yes, we can, but it didn't really take off until the phone came along . . .

Thus, we start with the advent of telephony.

An overarching view of telephony

Telecommunication didn't originate with the phone or phone-based communications systems. In fact, the term generally applies to any assisted transmission of signals at a distance. Smoke signals rank among its better known instances. However, the telecommunications networks to which we often refer also imply *telephony*.

A telecommunication system consists of three basic elements:

- ✓ **A transmitter:** Converts information to a signal
- ✓ **A transmission medium:** Carrier network for the signal
- ✓ **A receiver:** Converts a signal back into information

The *public switched telephone network* (PSTN) is that ordinary, old-fashioned telephone system (the transmission medium) you've come to know quite well. As a general term, it refers to the variety of worldwide telephone networks and services that support global analog-based communications. You may also hear it referenced as the plain old telephone service or POTS.

Circuit-switching telecommunication networks all utilize established fixed bandwidth circuits or channels between nodes and terminals. Users may communicate with one another through these devices as if nodes were physically connected to a single electrical circuit. Although circuit-switching can be used for more than voice traffic, this is where analog transmission begins.

Signaling systems eventually transitioned from analog to digital. Analog signals vary continuously across a definite signal range with respect to the transmitted information. A digital signal encodes information as a set of discrete values. During transmission, analog signals are susceptible to subtle noise degradation whereas digital signals remain intact until signal-to-noise levels exceed a certain threshold.

In analog telephone networking, callers are connected to each other by switches at various telephone exchanges, which form an electrical connection between both end-user points (the transmission medium). Switches are set electronically when callers dial the intended phone number, and once connected, a caller's voice transmission is transformed into an electrical signal through the caller's handset (the transmitter). This signal is transmitted electrically until received on the opposite end where it's then transformed back into sound by a speaker in the recipient's handset (the receiver).

Fixed-line telephones in most residences are analog, so that the speaker's modulated voice directly determines signal voltage levels. Local calls may be handled end-to-end as analog signals but telephone service providers typically convert signals to digital form for transmission before converting them back to analog for reception on the receiving end.

Wired communications get going

In 1831, Joseph Henry invented the first electric telegraph and four short years later Samuel Morse developed a coding system that took his name. In 1843, Samuel Morse then went on to develop the first long distance telegraph line and in that same year Alexander Bain patented the first fax machine. It wasn't until 18 years later that the Pony Express service commenced in the United States.

Within one adventurous century we saw a powerful communications culture spring to life, which concluded with patenting direct dial telephones (1889) and improvements to nascent wireless technology (1894). It would become the catalyst that would forever change the landscape of man-made communications.

Emerging data transmission

Computer networking as we know it today was born from the pet projects of the past. In 1940, George Stibitz successfully transmitted math problems from Dartmouth College in New Hampshire to his Complex Number Calculator in New York using a teletype system. This central server and dumb terminal concept flourished throughout the 1950s and remains in use today in the form of server-based cloud computing.

It wasn't until the 1960s that academic and government researchers began investigating packet-switching technology that allowed chunks of data to be sent between different computers. Such an arrangement obviates the need for a central server or mainframe and enables multi-node network transmission, which first emerged in 1969. Known as the Advanced Research Projects Agency Network (ARPANET), this network grew from four nodes in 1969 to 213 independent network nodes by 1981. Today, it's known as "The Internet" with millions of nodes online.

Virtual circuits and datagram packetry

Virtual circuit telecommunication is a connection-oriented service that makes delivery using packets. A packet-switching network splits data traffic — the digital representations of text, sound, and video — into discrete, well-defined chunks called *packets*. Each packet gets routed through a shared network labeled with its own destination and serialization identifiers. Packet-switching precludes the need for a dedicated circuit path to assist each transmission with delivery to its intended recipients.

Packet-switching contrasts principally with circuit-switching, which establishes a specific circuit and constant bit-rate between sender and recipient. Packet-based services may be delivered through virtual circuits using connectionless messages called *datagrams*. Virtual circuit networking operates in numbered layers, which designates the level of involvement for a given connection medium or protocol. Most commonly, these references refer to the seven-layer Open Systems

Interconnection (OSI) Network Reference Model (introduced in Chapter 1).

Virtual circuit layering

Under the OSI's Network Reference Model, each layer corresponds to a different protocol designation or family of designations. Layer 1 is the lowest layer, the Physical layer (covered in detail in Chapter 1), and defines electrical and physical specifications for devices. Most important, it defines the relationship between network devices and physical media including pin layouts, cable and voltage specifications, and network adapters.

Layer 2 virtual circuits are established at the Data Link layer (see Chapter 1), which provides functional and procedural means for transferring data between network entities. This layer was originally intended for point-to-point and point-to-multipoint media, as is characteristic of the telephone system.

Layer 3 of the OSI Basic Reference Model is the Network layer (again, see Chapter 1). This layer provides the functional and procedural means for transferring variable-length data sequences from source to destination using one or several network pathways. The best-known protocol at this layer is the Internet Protocol (IP), which manages the connectionless transfer of data one hop at a time from system to system, system to router, and so forth.



There are indeed four other layers residing above these, but layers one, two, and three are the primary focus of our discussion.

The Internet is born

ARPANET preceded the popular Internet infrastructure, a worldwide network of computers and computer networks that communicate using its namesake — the Internet Protocol. Each computer has a unique IP address that serves a similar purpose as the telephone exchange phone number for routing information to intended end-point recipients. As of 2008, Internet WorldStats.com estimates that 21.9 percent of the global

population has access to the Internet, with 73.6 percent of North Americans connected.

Data-bearing communications has forever transformed the landscape of man-made communications. What we're now witnessing is a potent combination of voice and data services from common sources.

Voice and data converge

Most telcos now function as Internet Service Providers (ISPs) and the distinction between these two separate entities has blurred with time. The trend for supplier convergence throughout the communications industry continues to this day.



Supplier convergence describes an increasing tendency for companies to offer combinations of services and products previously provided through separate entities.

The Internet's popularity derives in large part from its ability to support multiple protocols for voice and data applications. Although we're most familiar with using the Internet as a text, picture, and multimedia infrastructure, it is also quite capable of conveying rich media such as Voice over IP (VoIP) or even live video.

Transparent optical networking

Optical transport networking (OTN) — sometimes called G.709 or digital wrapper — is a next-generation method to pass information through optical networks. The key to understanding OTN is transparency. Transparent services are becoming increasingly popular for service providers as end-user networks incorporate media-rich applications and protocols. Transparency also extends to combining synchronous and asynchronous services on the same wavelength.

OTN's primary benefit is the export of transparent services. As end-users continue to deploy more complex networks, intermediary routing devices need access to SONET/SDH overhead bytes. A traditional SONET/SDH device strips these bytes and thereby limits its own functionality. OTN enables full manageability under these circumstances.

It's only fitting that voice and data signals combine into a cohesive unit for delivery. As lower-level layers — one, two, and three — blend together, the method of network transport achieves superior levels of performance. Along with such convergence also comes a unified view of the network.

A View of the Unified Network

Service providers continue to feel pressure to increase capacity and decrease costs. Deploying new optical transport equipment can lower capital expenses while simultaneously simplifying the network. The ultimate goal is to achieve true end-to-end control over a convergent packet-based and optical network that delivers a wider range of high-value services. Service providers must also provide backward compatibility while adding enhancements to networking services.

Network growing pains

Service providers have invested billions in the existing network infrastructure to support legacy applications, devices, and services — so there's no compelling reason to abandon it. Instead, it makes strong financial sense for service providers to extend existing infrastructure rather than switching to an entirely new one. IP-based-services growth remains steady despite gradual declines in per-bit-services revenues.

Lowering costs on wired subscriber services while stabilizing pricing for packet services is an important challenge. Enterprise subscribers also face the challenge of maintaining line costs while supporting emerging services.

New approaches to optical networking

Lowering capital and operating expenses to support legacy applications while moving forward with future planning and projects requires a new approach to optical networking. Service providers must learn how to deliver differentiated high-value service level agreement (SLA)-based services.

The next generation of network architectures and products can help manage the expense and risks involved in network convergence. This is possible because of:

- ✓ **Improved economies of scale:** New optical equipment facilitates construction of large-scale, multi-service networks that spread infrastructure costs across multiple services. Service providers need no longer build overlay networks.
- ✓ **Lowered capital expenses:** Higher line speeds, multiple wavelengths, improved switch density, service standardization, and programmable optics provide lower costs-per-bit than ever before.
- ✓ **Converging layers:** Combining protocols and services simplifies the network. The photonic, TDM, and data layers converge so that fewer single-purpose network elements are needed.
- ✓ **Simplifying operations:** A small percentage of network costs are capital; the remainder are operational. Low-touch network intelligence and control reduce new service labor costs.

Network layers are collapsing thanks to technology integration. Remember what we said about the lower layers of the OSI reference model merging? New WDM topologies — such as sub-wavelength switching and grooming, optical bypass, and reconfigurability — can support high-bandwidth packet video, voice, and data services effectively and efficiently.

Functions that help operators simplify network operations and reduce turn-up time include automated topology discovery, comprehensive end-to-end connection control, and management at both customer and service levels. Remote provisioning, remote reconfigurability, and vendor-independent interoperation also help lower operational expenses.

The proliferation of packet-based services and optimized network architectures also continue to improve network efficiency and further consolidate network elements. This benefits service providers and enterprise subscribers alike.



Earlier we stated that circuit-switching and packet-switching networks are separate entities — mutually exclusive media designed for different uses. Carriers are faced with new opportunities to deliver bundled high-value, high-margin services cost-effectively while managing the risks associated with migrating from circuit-switched services to packet-based infrastructures.

Ten Prognostications on Optical Networking

Here's a list of ten predictions we distilled from our preceding looks into the past, present, and future of optical networking. Buckle your seatbelts; this will go FAST!!!

- ✓ **Faster, faster, faster:** As optical technology innovations emerge, fiber bandwidth keeps increasing.
- ✓ **Cheaper bits:** More bandwidth, more coverage, and more services mean per-bit transport costs keep going down.
- ✓ **Rich services:** As bandwidth becomes more broadly available, clever engineers and software makers will come up with cool new things to do with it.
- ✓ **From triple-play to multi-play:** Today, cable companies, telcos, and ISPs all want you to get voice, data, and video from them. That list can only get longer over time.
- ✓ **Management and grooming rule the backbone:** Bigger bandwidth cries out for more and better ways to aggregate and manage big, fat optical data pipes. Those who innovate best will profit most.
- ✓ **Roll your own infrastructure:** Organizations and enterprises that bite the bullet and build their own optical infrastructures will realize the biggest benefits from innovation. Dark fiber combined with their own equipment is an unbeatable combination.
- ✓ **Unified networking rules the world:** Today, unified networking is still in its infancy. Rich applications are just barely able to combine e-mail, voicemail, documents, and data. As unification proceeds, capabilities get a lot more interesting.

- ✓ **Do more with the same fiber:** Improvements in optical receivers and emitters will let companies and organizations keep doing more with the same optical fibers. The sky may not even be the limit!
- ✓ **Here, there, and everywhere:** More, better, and faster network connections will continue to erase the impact of distance and location on collaboration and interaction. We already have telemedicine and teleworking; get ready for lots more.
- ✓ **There's always TNBT (the next big thing):** Ever-increasing backbone and user-interface speeds will provide the bandwidth to accommodate the “next big things” — such as grid computing. Transport and data (Ethernet) speeds are converging at 10 Gbps, and 40 Gbps is finally seeing deployment in the WAN backbone. But, the insatiable appetite for bandwidth will quickly eat up 40 Gbps capacity. Recent demonstrations of 100 Gbps in data networks as well as field tests by carriers point toward a future backbone capable of transmitting 20 million 4-drawer file cabinets full of data within 12 hours. And on it goes . . .

TimBuk3 said “The future’s so bright, you gotta wear shades.” Although they didn’t know it when they wrote it, they were talking about optical networking.

A Glossary of Acronyms

.....

1000BaseLX: an optical Gigabit Ethernet standard
ADM: add-drop multiplexer
APD: avalanche photodiode
ASON: Automatically Switched Optical Network
ATM: Asynchronous Transfer Mode
bps: bits per second
CAT5: Category 5, a type of twisted pair networking cable
C-band: Optical spectrum (color band) between 1530 and 1565 nm. Not to be confused with the microwave C-Band.
CBR: constant bit rate
CRC: cyclic redundancy check
CWDM: Coarse Wave Division Multiplexing
DFB: Distributed Feedback laser
DS1: Digital Signal 1
DS3: Digital Signal 3
DWDM: Dense Wave Division Multiplexing
E-1: European Carrier Level 1
E2000: fiber optic connector name, European design
EDFA: Erbium Doped Fiber Amplifier
E-LAN: Ethernet virtual local area network
EMI: electro-magnetic interference
EML (or EAM): externally-modulated laser or electro-absorption modulator-laser
E-NNI: External Network-to-Network Interface
EPL: Ethernet Private Line
EVC: Ethernet Virtual Connection
EVPL: Ethernet Virtual Private Line
FC: Ferrule Connector or Fiber Channel. Threaded optical connector.
FEC: Forward Error Correction
FP laser: Fabry Perot laser
FR: Frame Relay
G.709: ITU-T Recommendation "Interfaces for the Optical Transport Network"
GbE: Gigabit Ethernet (10 GbE = ten gigabit Ethernet, 100 GbE = hundred gigabit Ethernet)
Gbps: Gigabits per second
GMPLS: Generalized Multi-protocol Label Switching
IETF: Internet Engineering Task Force
I-NNI: Internal Network-to-Network Interface
IP: Internet Protocol
ISP: Internet Service Provider
ITU: International Telecommunications Union

Km: kilometer

LAN: Local Area Network

Laser: light amplification by stimulated emission of radiation

L-Band: Optical spectrum between 1575 and 1610nm. Not to be confused with the microwave L-band.

LC: Lucent Connector. Miniature push-on optical connector.

LTE: line terminating equipment

MAC: Media Access Control

MAN: Metropolitan Area Network

Mbps: Megabits per second

MHz: MegaHertz (millions of cycles per second)

MMF: multi-mode fiber

MPLS: Multi-protocol Label Switching

MSPP: Multi-Service Provisioning Platform

MT-RJ: Mechanical Transfer Registered Jack

NE: Network Element

Nm: nanometer, 1-billionth of a meter

NMS: Network Management System

OADM: Optical Add-Drop Multiplexer

OC-n: Optical Carrier Level n (1, 3, 12, 48, 192, 768 mentioned)

OEO (sometimes O-E-O): optical-electronic-optical conversion

OLA: Optical Line Amplifier

OSI: Open Systems Interconnection

OTN: optical transport networking (see G.709)

O-UNI: Optical User-network Interface

OXC: Optical Cross-Connect

PIN: a semiconductor made with a layer of P-type material, an integral silicon layer, and a layer of N-type material

PMD: polarization mode dispersion

POTS: Plain Old Telephone System

PSTN: Public Switched Telephone Network

QoS: Quality of Service

ROADM: Reconfigurable optical add-drop multiplexer

SC: Subscriber Connector or Standard Connector Push-on optical connector

SDH: Synchronous Digital Hierarchy

SLA: Service Level Agreement

SMF: single-mode fiber

SONET: Synchronous Optical Network

ST: Straight Tip. $\frac{1}{4}$ turn locking optical connector.

Tbps: Terabits per second

TCP: Transmission Control Protocol

TCP/IP: Transmission Control Protocol/Internet Protocol

TDM: Time Division Multiplexing

TMF: Telemanagement Forum

VCSEL: Vertical Cavity Surface Emitting Laser

VLAN: Virtual Local Area Network

WAN: Wide Area Network

WDM: Wave Division Multiplexing



**Understand the benefits
of optical networking**

Make the most of this high-speed technology

Are you faced with budgetary or implementation decisions for networking infrastructure? Want to avoid expensive carrier tariffs for wide-area networking? With this book, you not only get the basics of optical networking, but you also find out how to save big bucks and blow the lid off your wide-area network capacity. You also see that when you build your own optical network to convert service costs to capital expenses, you open up that network to future growth with the purchase of equipment upgrades. Along the way, you get acquainted with the various types of optical equipment, what kinds of distances they can cover, and how you can put them to work for your company or organization.

**THE
DUMMIES
WAY®**

Explanations in plain English
"Get in, get out" information
Icons and other navigational aids
Top ten lists
A dash of humor and fun

Discover how to:

*Make optical net-
working work for you*

*Choose optical net-
working technologies*

Light up "dark fiber"

*Reduce recurring
costs*

*Bust loose from
expensive carriers
or telcos*

**Get
smart!**
www.dummies.com

- ✓ Find listings of all our books
- ✓ Choose from many different subject categories
- ✓ Sign up for eTips at etips.dummies.com