

Routing **first-step**

Your first step into the world
of **routing**

- No routing experience required
- Includes clear and easily understood explanations
- Makes learning easy



Routing **first-step**

Bill Parkhurst



Cisco Press
800 East 96th Street
Indianapolis, IN 46240



Routing first-step

Bill Parkhurst

Copyright© 2005 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

First Printing September 2004

Library of Congress Cataloging-in-Publication Number:

2003116566

ISBN: 1587201224

Warning and Disclaimer

This book is designed to provide information about routing. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The author, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Publisher

John Wait

Editor-in-Chief

John Kane

Cisco Representative

Anthony Wolfenden

Cisco Press

Program Manager

Nannette M. Noble

Production Manager

Patrick Kanouse

Development Editor

Dayna Isley

Senior Project Editor

San Dee Phillips

Copy Editor

Laura Williams

Technical Editors

Mark Gallo

Tyler Hodges

Kevin Turek

Editorial Assistant

Tammi Barnett

Book and

Cover Designer

Louisa Adair

Composition

Mark Shirar

Indexer

Tim Wright

Proofreader

Karen A. Gill

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Corporate and Government Sales

Cisco Press offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales.

For more information please contact: U.S. Corporate and Government Sales 1-800-382-3419
corpsales@pearsontechgroup.com

For sales outside the U.S. please contact: International Sales international@pearsoned.com

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.



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

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

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

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco.com Web site at www.cisco.com/go/offices.

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

Copyright © 2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco Powered Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, IQ Net Readiness Scorecard, Networking Academy, and ScriptShare are trademarks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, Gigabit, Internet Quotient, IOS, IPTV, IQ Expertise, the IQ logo, LightStream, MGX, MiCA, the Networkers logo, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, Ram-MUX, Registrar, SlideCast, SMARTnet, StrataView Plus, Stratus, SwitchProbe, TeleRouter, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0303R)

Printed in the USA



About the Author

Bill Parkhurst, Ph.D., CCIE No. 2969, is an advisory engineer at Cisco Systems and a technical advisor for Cisco certifications. Bill has taught networking to students on five continents at all levels, from beginner to advanced. Before joining Cisco, Bill taught networking at Wichita State University and now has the pleasure of working with many of his former students.

About the Technical Reviewers

Mark Gallo is a technical manager with America Online where he leads a group of engineers responsible for the design and deployment of the domestic corporate intranet. His network certifications include CCNP and CCDP. He has led several engineering groups responsible for designing and implementing enterprise LANs and international IP networks. He has a bachelor of science degree in electrical engineering from the University of Pittsburgh. Mark resides in northern Virginia with his wife, Betsy, and son, Paul.

Kevin Turek, CCIE No. 7284, joined Cisco in 2000 and is currently working as a network consulting engineer in the Cisco Federal Support Program in Research Triangle Park. He has been involved with several Cisco Press projects, including co-authoring the Cisco Press title *Cisco Catalyst QoS: Quality of Service in Campus Networks*. Kevin earned his bachelor of science degree in business administration at the State University of New York, Stony Brook.

Dedications

This book is dedicated to my daughter, Laura. I want to congratulate her on graduating from Kansas State University and taking her own first step. Now, if I can only get her to read this book.



Acknowledgments

I would like to acknowledge the effort and support of all those involved with the development of this book. From Cisco Press, I want to thank John Kane for having the confidence in me to suggest that I write this book; and Dayna Isley for her unmatched efforts in the development process, and for never letting me take the easy road. The technical reviewers, Mark Gallo and Kevin Turek, for diplomatically finding and pointing out errors, and for their insight regarding the content and flow. I want to especially thank my friend Tyler Hodges for reviewing this book from the perspective of a young and open mind. Tyler's comments were extremely helpful in ensuring this book was appropriate for the intended audience, and that everything was explained clearly. Finally, I want to thank my wife, Debbie, for her support during the writing of this book, and for the initial proof-reading she did to make sure I appeared literate to my editors and reviewers.

Contents at a Glance

	Introduction	xiv
Chapter 1	Routing and Switching in Everyday Life	3
Chapter 2	A_{16} B_{16} C_{16} , As Easy As 01_2 , 10_2 , 11_2	39
Chapter 3	Internet Addressing and Routing	63
Chapter 4	Routing IP	123
Chapter 5	Cisco Interior Gateway Protocols	165
Chapter 6	Open Shortest Path First—Better, Stronger, Faster	195
Chapter 7	Intermediate System-to-Intermediate System—Better, Stronger, Faster, and Scarier	245
Chapter 8	Border Gateway Protocol—The Glue That Holds the Internet Together	279
Chapter 9	Multicast—What the Post Office Can't Do	337
Appendix A	Answers to Chapter Review Questions	367
Glossary		393
Index		403



Contents

Introduction xiv

Chapter 1 Routing and Switching in Everyday Life 3

Postal System 3

Package Protocol 6

Addressing Protocol 6

Mail Delivery Protocol 8

Local Delivery 9

Local Off-Street Delivery 15

Citywide Delivery 18

Statewide Delivery 21

Countrywide Delivery 23

Highway System 29

Telephone System 33

Summary 36

Chapter Review Questions 36

Chapter 2 A_{16} B_{16} C_{16} , As Easy As 01_2 , 10_2 , 11_2 39

Decimal Numbering System 39

Binary Numbering System 40

Octal Numbering System 44

Hexadecimal Numbering System 45

Dotted Decimal Notation 47

Conversions Between Number Systems 47

Binary to Octal 48

Binary to Hexadecimal 48

Binary to Decimal 49

Binary to Dotted Decimal 49

Octal to Binary 50

Octal to Hexadecimal 51

Octal to Decimal 51

Octal to Dotted Decimal 51

Hexadecimal to Binary 52

Hexadecimal to Octal 52

Hexadecimal to Decimal	53
Hexadecimal to Dotted Decimal	54
Decimal to Binary	54
Decimal to Octal	56
Decimal to Hexadecimal	56
Fun Binary Number Facts	57
Summary	60
Chapter Review Questions	61

Chapter 3 Internet Addressing and Routing 63

Internet Addressing	63
Internet Addressing Protocol	68
Classful IP Addresses	72
Private IP Addresses	74
Address Resolution	75
Intra-LAN Communication	77
Inter-LAN Communication	81
IP Header Format	85
TCP/IP Layered Protocol Model	90
Classless Internet Addressing	94
IP Routing and Route Summarization	105
Supernets	113
IP Version 4 and IP Version 6	113
IPv6 Address Format	116
IPv6 Address Types	117
Summary	118
Chapter Review Questions	118
References	119

Chapter 4 Routing IP 123

Delivering Snail Mail and E-Mail—Any Difference?	124
Basic IP Router Configuration	125
Routing Information Protocol Version 1	130
Counting to Infinity	136
Split Horizon	137
Poison Reverse	138



Convergence	139
Variable-Length Subnet Masks	142
Hop Count Limitation	146
RIPv1 Algorithm	147
RIP Version 2	150
Security	153
Route Summarization	155
Summary	160
Chapter Review Questions	161
References	163

Chapter 5 Cisco Interior Gateway Protocols 165

Introducing IGRP	165
IGRP Metrics	168
IGRP Limitations	171
Introducing EIGRP	174
EIGRP Neighbor Discovery	174
Basic EIGRP Configuration	175
EIGRP Metrics	178
EIGRP Route Update Algorithm	179
EIGRP Network Summarization	186
Comparing IGRP and EIGRP	189
Summary	190
Chapter Review Questions	191
References	192

Chapter 6 Open Shortest Path First—Better, Stronger, Faster 195

OSPF Areas	196
Link States	198
OSPF Router ID	201
Basic OSPF Configuration	203
OSPF Neighbor Discovery	209
OSPF Timers	212
OSPF Metrics	213
OSPF Router Types	214

OSPF Route Types	217
OSPF Area Types	219
Stub Area	220
Totally Stubby Area	221
Not-So-Stubby Area (NSSA)	222
Totally Not-So-Stubby Area	226
OSPF External Route Summarization	227
OSPF Route Summarization	229
OSPF Virtual Links	230
Selecting the Shortest Path	234
OSPF LSA Types	238
Summary	239
Chapter Review Questions	240
References	242

Chapter 7	Intermediate System-to-Intermediate System—Better, Stronger, Faster, and Scarier	245
	Comparing IS-IS and IP Networks	245
	IS-IS Areas	249
	IS-IS Link States	254
	Basic Single Area IS-IS Configuration	256
	IS-IS Metrics	261
	IS-IS Multiple Area Configuration	264
	IS-IS Route Summarization	270
	Route Leaking	272
	Comparing IS-IS and OSPF	273
	Summary	274
	Chapter Review Questions	275
	References	276

Chapter 8	Border Gateway Protocol—The Glue That Holds the Internet Together	279
	Understanding the Need for BGP	280
	BGP Attributes	282
	Autonomous System Path Attribute—AS_PATH	284
	WEIGHT Attribute	287



Local Preference Attribute—LOCAL_PREF	288
Metric or MULTI_EXIT_DISC (MED) Attribute	289
ORIGIN Attribute	290
NEXT_HOP Attribute	291
Community Attribute	292
Basic BGP Configuration—EBGP	294
Advertising IP Prefixes	298
Using the network Command to Inject Routes	299
Using Redistribution to Inject Routes	305
Using BGP to Inject Routes	307
Basic BGP Configuration—IBGP	309
IBGP and Loopback Interfaces	320
Scaling IBGP	323
Route Reflector	324
Confederations	325
BGP Route Summarization	328
BGP Decision Process	332
Summary	333
Chapter Review Questions	333
References	334

Chapter 9	Multicast—What the Post Office Can't Do	337
	Comparing Unicast and Multicast Routing	337
	Multicast Switching	340
	IGMP	343
	CGMP	345
	IGMP Snooping	346
	Multicast Forwarding	346
	Protocol Independent Multicast Routing Protocol	349
	PIM Dense Mode	349
	PIM Sparse Mode	353
	Static RP	355
	Auto-RP	356
	PIM SM Version 2	357
	Anycast RP	359

Reserved Multicast Addresses 362

Summary 362

Chapter Review Questions 363

References 364

Appendix A Answers to Chapter Review Questions 367

Glossary 393

Index 403



Introduction

Over the last decade, the Internet has grown from an interesting research project to a ubiquitous form of communication that has forever changed our world. E-mail, instant messaging, IP telephony, music and video on-demand, online banking, gaming, and travel planning are but a few of the many applications that have made many of our lives more convenient. In time, the Internet will eventually be available to everyone—and we can only guess what applications are to come. But whatever form the Internet takes, and whatever applications become available, there is an ever-increasing need for people to design, deploy, support, manage, teach, and sell these technologies.

The Cisco Press First-Step series is the starting point for understanding the basics of computer networking, and this book covers the concepts and protocols that enable the transfer of information from anywhere to anywhere using computer networks. No prior knowledge is assumed and each new concept is explained in understandable terms, allowing you to grasp the concepts without getting lost in the lingo. Of course, you will learn some of the lingo of the trade, but I assure you that it will be painless (mostly). The following description of the chapters in this book will give you an idea of what is in store for you:

- Chapter 1, “Routing and Switching in Everyday Life”—Examines familiar systems, such as postal, highway, and phone, that route information similar to the Internet. These systems are used to present an understanding of the concepts and terminology of routing and switching.
- Chapter 2, “ A_{16} B_{16} C_{16} , As Easy As 01_2 , 10_2 , 11_2 ”—Covers number systems used by networking professionals. If you like numbers, this chapter should be fun. If you don’t like numbers, hopefully this chapter will change your mind.
- Chapter 3, “Internet Addressing and Routing”—Uses the concepts learned in the first two chapters to discuss the addressing scheme used in computer networks in detail.

- Chapter 4, “Routing IP”—Introduces network routing, basic configuration of Cisco routers, and RIP, one of the first network routing protocols.
- Chapter 5, “Cisco Interior Gateway Protocols”—Describes the basic concepts and configuration of the Cisco routing protocols: IGRP and EIGRP.
- Chapter 6, “Open Shortest Path First—Better, Stronger, Faster”—Introduces OSPF, a popular standards-based routing protocol used in many networks. After reading this chapter, you will understand the concepts, operation, and configuration of OSPF.
- Chapter 7, “Intermediate System-to-Intermediate System—Better, Stronger, Faster, and Scarier”—Describes IS-IS, which is similar to OSPF but has a few differences that tend to scare people away from using it. This chapter will remove the mystery and darkness that surrounds IS-IS.
- Chapter 8, “Border Gateway Protocol—The Glue That Holds the Internet Together”—The routing protocols covered in Chapters 4 through 7 are used in the networks of companies, universities, government agencies, and so on. The Internet is what ties all these networks together, and BGP is the protocol that holds it all together.
- Chapter 9, “Multicast—What the Post Office Can’t Do”—Presents the concepts and routing protocols used to simultaneously send information to more than one recipient, something the post office can’t do.
- Appendix A, “Answers to Chapter Review Questions”—Each chapter ends with a “Chapter Review Questions” section that helps reinforce your understanding of the topics discussed. This appendix repeats the questions and provides the answers.
- Glossary—Throughout this book many new terms will be introduced and defined. These new terms are all listed in the glossary for easy reference.



Who Should Read This Book

This book, and this series, is for anyone interested in starting down the road of learning about the many facets of computer networking. The First-Step series can give you a solid foundation to forge ahead into any area of networking you find interesting. If you are in high school and are exploring possible career goals, if you are a college student and want to see if this field is something you might be interested in pursuing, if you are already working in another hi-tech field and are considering a career change, or if your company depends on computer networks but you only need to understand the basics for making the right business decisions, this book is for you.

If you are interested in checking out other books in the First-Step series, visit <http://www.ciscopress.com/firststep> for more information.

What You Will Find in This Book

This book includes several features to help ease the process of learning about routing:

- **Chapter objectives**—Every chapter begins with a list of objectives that are addressed in the chapter. The objectives are revisited in the chapter summary.
- **Highlighted key terms and Glossary**—Throughout this book, you will see terms formatted with bold and italic. These terms are particularly significant in routing. If you find that you aren't familiar with the term or that you need a refresher, look up the term in the Glossary toward the end of the book to find a full definition.
- **Chapter summaries**—Every chapter concludes with a comprehensive summary that reviews chapter objectives, ensuring complete coverage.
- **Chapter review questions**—Every chapter concludes with review questions. These questions test the basic ideas and concepts covered in each chapter. You can find the answers to the questions in Appendix A.

- **Examples**—In networking, you will encounter router configurations and output. To help prepare you, this book provides basic configuration and output examples that are thoroughly explained in the text and, when necessary, annotated.

The illustrations in this book use the following icons for networking devices:



PC



Router



Switch



Network Cloud

What You Will Learn

After reading this chapter, you should be able to

- ✓ Understand the concepts of routing and switching using familiar systems (postal, road, and telephone)
- ✓ Understand the properties of a routable address and how various parts of the address are used to route something from point A to point B
- ✓ Understand the concept of route summarization, or aggregation, at the core and distribution level for reducing the amount of information required in the routing tables at these levels



CHAPTER 1

Routing and Switching in Everyday Life

The Internet is nothing new. Think about that statement. Certainly the technologies are new and applications such as e-mail, instant messaging, web sites, and video on demand are new, but the concepts are not. How the Internet operates is conceptually no different than how the postal, highway, and telephone systems operate. They all perform the same function—getting things from one place to another. The postal system does this with mail, the highway system does this with families on vacation and produce from California, and the phone system does this with phone calls. This book is about getting things from here to there. If you understand the concepts behind the postal, highway, and phone systems, you can understand the concepts of computer networks. After the concepts are mastered, learning the details of how it all works is easy. So before we jump off into the amazing world of the Internet, let's take the time to look at systems you are already familiar with.

Postal System

First, let's take an imaginary trip back in time and pretend it's your job to develop a system for delivering letters and packages, and whatever else people can think of sending, to anyone anywhere in the country. So, where do you start? I think the best place to start is to establish rules. If your system has any chance of working, you need everyone to follow a clear set of rules. You also want to give your set of rules a fancy name so when you present your plan you sound somewhat authoritative.



Consulting your pocket dictionary, you find that a set of rules is also a set of conventions that is referred to as a *protocol*. That sounds fancy enough, so let's call your set of rules a protocol.

How many protocols do you need? That depends on how you structure the mail delivery system, so you need to think about how the system operates. Ultimately, you need to deliver something from point A to point B. These points can be located anywhere in the country.

After some serious thought, you decide you need to develop the following protocols:

- **Package Protocol (PP)**—The type and parameters (size, weight, contents, and so on) of the mail that can be delivered.
- **Addressing Protocol (AP)**—Where is the item coming from and where is it going to.
- **Delivery Protocol (DP)**—What is the most efficient way to deliver an item from point A to point B.
- **Transportation Protocol (TP)**—How to physically deliver the mail to the proper recipient.

These are not protocols used by the post office, of course, but are used for the purpose of the discussion regarding the routing of mail.

In addition, the protocols should be as independent as possible and scalable. Independence means that the operation of one protocol should not depend on other protocols. For example, the type of mail sent should not affect the addressing protocol. You don't want to use one addressing protocol for letters and another for packages. If you send a letter to your mother and a birthday present to your sister, you would have to use the same format for the address on both. For example, you might change the protocol specifying the location of the address on a letter, or package, or include additional information in the address (a ZIP code, for example). You should be able to do this without affecting which route the mail takes to its final destination.

The delivery protocol will make a delivery, or routing, decision based on the address. In that sense, the delivery protocol is dependent on the addressing protocol. You, however, want the freedom to change the addressing protocol without having to change the delivery protocol. Finally, the means that you use for the physical delivery of the mail (horse and wagon, truck, and so on) should not be dependent on what is actually being sent. With these constraints in mind, you can create a layered model called the Mail System Layered Protocol Model shown in Figure 1-1.

Figure 1-1 Mail System Layered Protocol Model

Package	What is being sent?
Address	Where is it going? Where is it from?
Delivery	What is the best route to use for delivery?
Transportation	How will the mail physically be delivered?

A *scalable protocol* means that the protocols still work well as the size of the system grows. For example, assume that mail will be delivered only between Chicago and Kansas City, and you will allow only one piece of mail every week. Because this is a small volume of mail, you hire someone with a horse to make the delivery. Assume this system works well, and you increase the volume to two pieces of mail per week. You could hire another person with a horse, but now you have two people to pay and two horses to feed. If you follow this method, you cannot scale your system to deliver thousands of pieces of mail because this would require thousands of people and thousands of horses! This system does not scale. You should require each rider to carry as much mail as possible so the system can scale to accommodate higher volumes of mail.

Now that you have decided which protocols are required, you need to determine the details of each protocol.



Package Protocol

The Package Protocol, or PP, is a relatively straightforward protocol. You have to decide what can be sent and the physical limitations of a piece of mail, such as size and weight.

Addressing Protocol

To deliver mail, you need to know where a particular piece of mail is going. In addition, it would be nice to know where a piece of mail came from in case it has to be returned. Therefore, you need a set of addressing rules that are part of our Addressing Protocol (AP). The two rules you need for the AP are the following:

- To Address rule that identifies where the mail is going
- From Address rule that identifies where the mail came from

But again, To Address and From Address are not sophisticated terms. Remember you do want your plan to sound sophisticated. Instead of To Address and From Address, use *destination address* and *source address* in your AP.

The addresses in the protocol need a defined structure. In other words, what information must be contained in a destination or source address? If an address contains multiple pieces of information, in what order should this information be presented? For the destination address, you need to know

- Who the mail is going to
- Where they live (state, city, street, and street number)

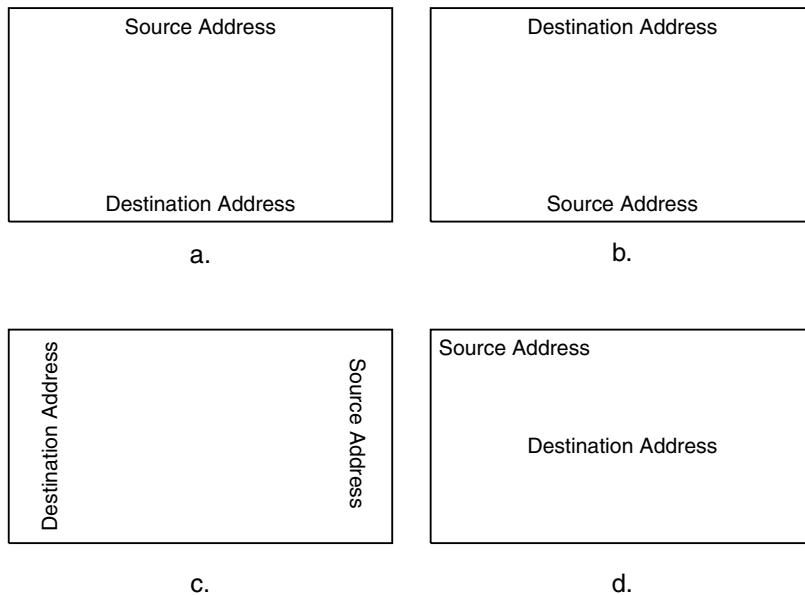
For the source address, you need to know

- Who the mail is from
- Where they live (state, city, street, and street number)

The addressing protocol needs to describe the exact location on the package for the source and destination addresses so they can easily be found. You also need to decide the placement order of the information contained in an address. After

experimenting with a number of address placement schemes, as shown in Figure 1-2, you decide to use the format shown in Figure 1-2 d. The beauty of this scheme is that it applies regardless of the type of letter or package sent through our mail system. Could the other schemes have been used? Yes. You only need to choose one and be consistent.

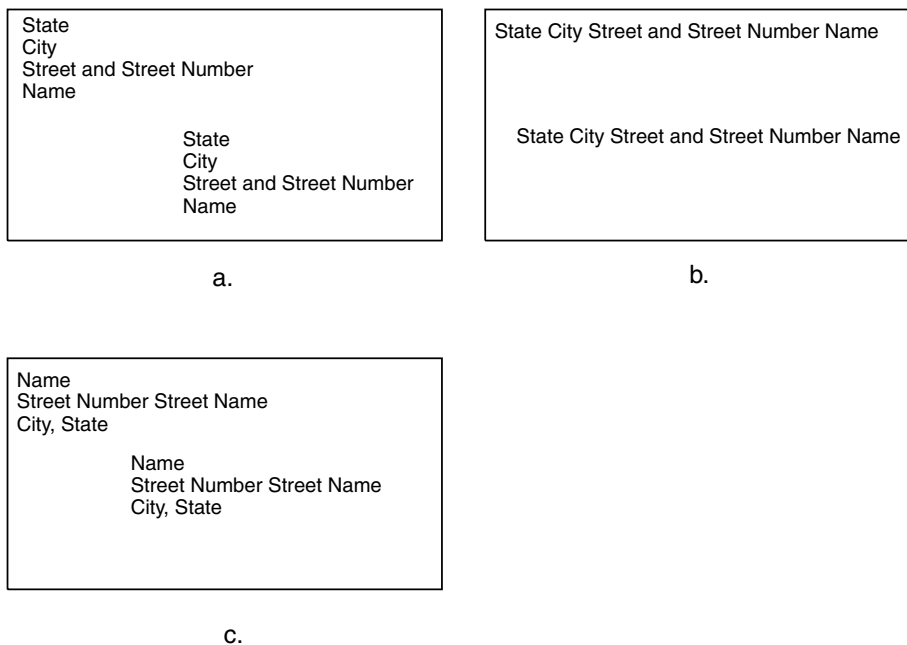
Figure 1-2 Rules for Source and Destination Address Placement



Various address information-ordering schemes were examined. (See Figure 1-3.) You have decided on the scheme in Figure 1-3 c. The other schemes would work just as well, of course, but you need to choose one and be consistent.



Figure 1-3 Address Ordering Schemes



The previous comment might seem obvious, but it is extremely important. Imagine a mail delivery system where the addressing information and placement are different for each type of letter, or package, that is delivered. By having an address-ordering scheme, it doesn't matter what you send—the addressing scheme is the same, independent of the type of letter or package.

Now that you have a workable addressing protocol, the next step is to develop a protocol for mail delivery.

Mail Delivery Protocol

The Mail Delivery Protocol, or MDP, is concerned with examining the information required by the AP and making a delivery decision. Mail will be delivered based on the destination address portion of the AP. The source address is needed if the mail needs to be returned. In that situation, the source address would become

the destination address. Let's look at the information contained in the destination address, and see what you can use to make a delivery decision. The destination address contains

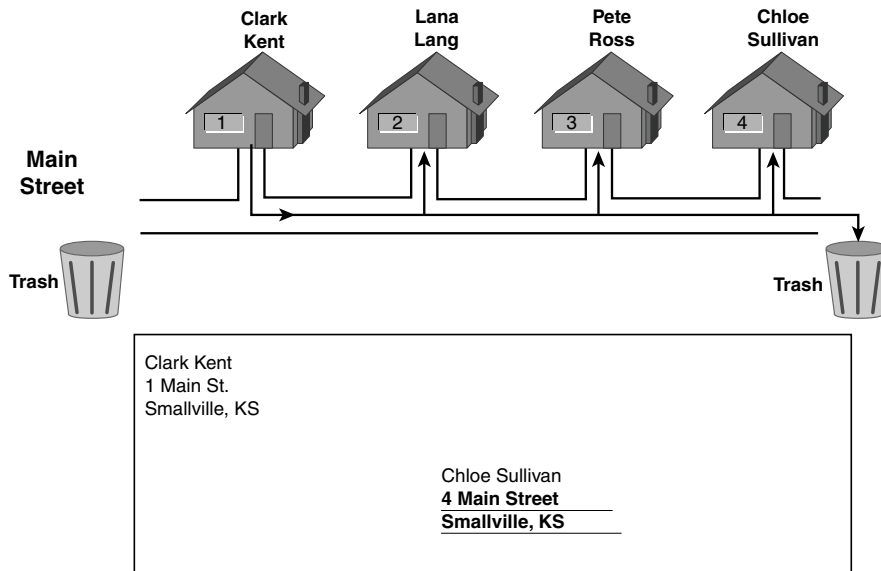
- Who the mail is going to
- Where they live (state, city, street, street number)

Further assume that all delivery facilities (post offices) will use all the information in the destination address to determine how to deliver the package. This being the case, every post office needs to maintain a list of every possible destination address—which would be every person in the country. This approach has many problems. The first problem is that each post office will need to maintain a big book that contains everyone's address. The second problem is that whenever anyone moves, or when someone immigrates to the United States, or emigrates from the United States, every address book at every post office would need to be updated. This scheme will simply not work. It is, in fact, unscalable. What you need is a MDP that will scale.

Go back and look at the information contained in the destination address and see if you can devise a delivery system that scales. You have already seen that maintaining a list of every person at every post office will not work. Can you devise a system that uses only parts of the destination address? The geographical relation between the source and the destination might provide us with ideas for setting up a scalable system for the delivery of mail at the local, city, state, and country levels.

Local Delivery

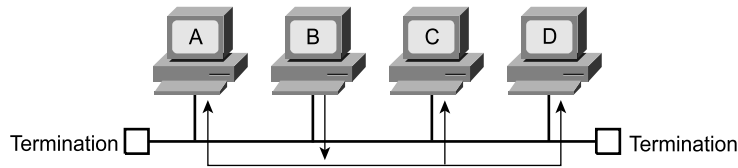
Local delivery is for mail with the same street name, city name, and state name in both the destination and source addresses. The additional information that is needed to deliver the mail is the house number. You do not need the information regarding to whom the package is being delivered. The MDP's job is over when the package arrives at the proper address. Consider two schemes for local mail delivery. The first scheme is called the Broadcast Local Delivery Protocol (BLDP) as shown in Figure 1-4.

Figure 1-4 Broadcast Local Delivery Protocol (BLDP)

In Figure 1-4, Clark sends a letter to Chloe. The mail carrier looks at all the information in the destination address except for the “who” component. If the street name is Main, the city name is Smallville, and the state name is Kansas, the letter carrier will take the letter to every house on Main Street and ask the resident if the letter is for them. If it is, the letter has been delivered. If not, the carrier moves on to the next house. If the letter has not been delivered after visiting the last house, the carrier will throw the letter away or return it to Clark.

What if the street name is not Main? Or the destination address has a different city or state name? At this point, you haven’t come up with a way to handle this situation, but you will in the following sections.

You’re probably wondering why you would even consider such a goofy way to deliver the mail. In the early days of computer networking, this is essentially how computers exchanged mail. As shown in Figure 1-5, a coaxial cable, like the type used for cable television, connected local computers.

Figure 1-5 Shared LAN

If computer B wanted to send a message to computer A, computer B would set the destination and source addresses to A and B, respectively. (Yes, computers have addresses. In fact, they typically have more than one address. We will get to that in Chapter 3, “Internet Addressing and Routing.”) The message was transmitted on the cable, and every computer connected to the cable received the message. Every computer had to take the time to examine the destination address and decide if the message was for them. If it wasn’t, the message was discarded. The terminations shown in Figure 1-5 serve the same purpose as the trashcans in Figure 1-4. They absorb the signal on the cable having the effect of throwing the signal away.

This was an inexpensive way to create a network of local computers, or a local-area network (LAN). All you had to do was buy a network card for each computer, run a cable from computer to computer, and connect the computers to the cable. It was inexpensive, but it had a number of annoying problems. If there was a break anywhere in the cable, the communication was lost between all computers. And if your cable was running between computers on multiple floors, it could take some time to find the break. The longer it took to find the break, the more upset the computer users became.

In addition, because the cable was a shared media, only one computer at a time could use the cable for communication. How did the computers know whose turn it was to use the cable? They didn’t. So they had to be able to detect a collision, which happens when two or more computers try to send something at the same time. As more people were connected to the network, the number of collisions increased, and the network became slower. Result? More upset computer users.



Ethernet was the name given to the predominant technology used for computer communication on a LAN. Before Einstein, scientists thought that the universe had to contain a substance through which light could travel. Otherwise, how could the light from distant stars reach Earth? This substance was given the name ether. The existence of ether was later disproved, but I guess someone involved with the early development of networks liked the term. Because computer communication had to travel through something, why not call it the ether? And because we are talking about networks, just use the term *Ethernet*.

Every Ethernet network card has an address called an Ethernet address. An example of an Ethernet address is 00-03-47-92-9C-6F. We make sense of that in Chapter 3.

Obviously, you do not want to use the previous method for local mail delivery. You would need a mail carrier for every street, and that does not scale. A better method is to use a local post office that has the responsibility for receiving and delivering mail to the local area. In Figure 1-6, Clark sends a letter to Pete. The letter is first sent to the local post office where the destination address is examined. Because the street, city, and state in the destination address is a street that this local post office serves, the letter is switched back to Pete's house on Main Street. The street names are the same, so the switch has to use only the house number to deliver the mail.

As technology improved, a better method for connecting computers on a LAN was developed, and it is similar to the local "switching" of mail. The switch in Figure 1-7 serves the same purpose as the local post office.

Figure 1-6 Switched Local Mail Delivery

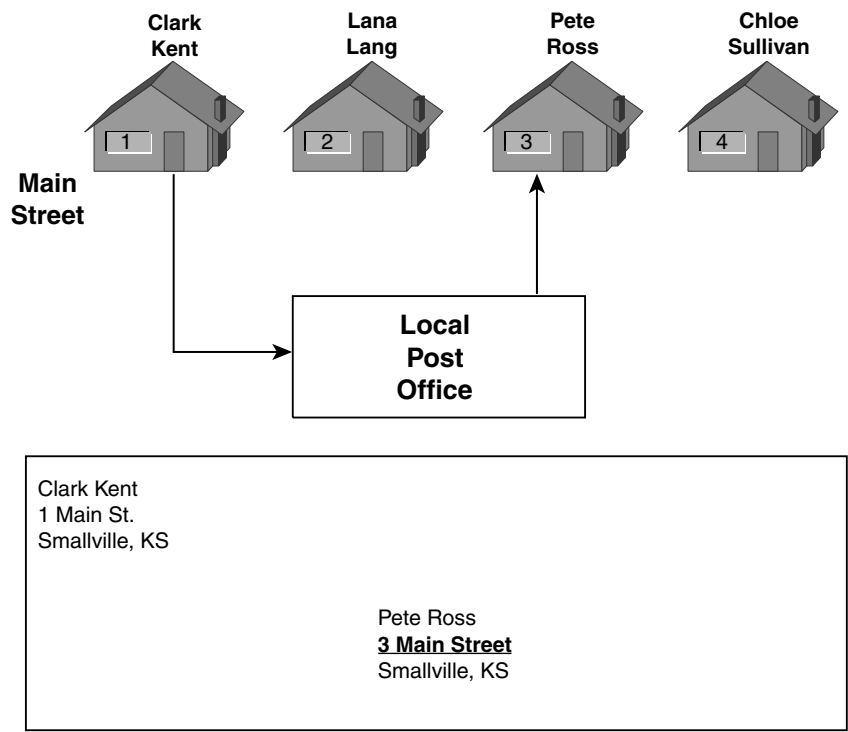
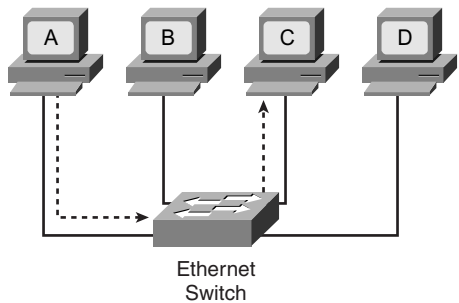


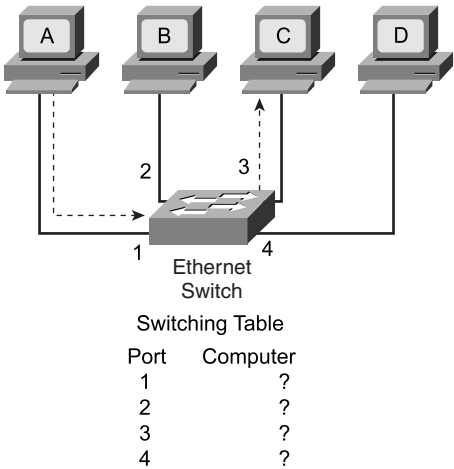
Figure 1-7 Switched LAN





When computer A wants to send a message to computer C, the message is sent to the Ethernet switch. The switch looks at the destination address and sends the message to computer C. Computers B and D never receive the message. How does the switch know which connection, or port, goes to computer C? Initially, it doesn't. The switch has to learn the location of every computer on the network, and build a switching table that contains a listing of each computer's address, and the port to which it is connected. When the switch in Figure 1-8 is first installed, the switching table is empty.

Figure 1-8 LAN Switching Table



When the switch receives a message from computer A, the switch learns that computer A is on port 1. How? The message was received on port 1, and the source address is A, so the switch knows which computer sent the message. The destination address is computer C, but the switch does not know to which port computer C is connected. In this case, the switch sends the message out all ports except for the one on which it was received. Therefore, computers B and D receive the message even though it was not addressed to them. Eventually, the switch learns on which port each computer is connected. Of course, that assumes that each computer

eventually sends a message. After the switching table is complete, the switch is able to forward the message to the computer for which it was meant. Table 1-1 contains the final switching table.

Table 1-1 Final Switching Table

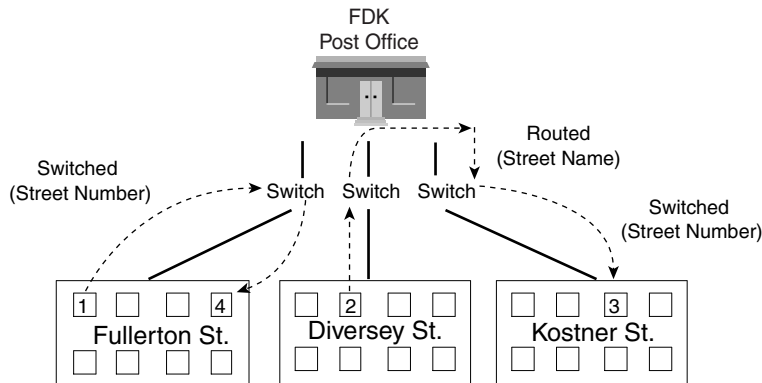
Port	Computer
1	A
2	B
3	C
4	D

Because the switching method of delivering mail seems to make the most sense, you adopt that as your local mail delivery method. Time to move on to local off-street delivery.

Local Off-Street Delivery

Local off-street delivery is for mail on different streets in the same city and state that can be accessed by the same post office. The information required to deliver the package is the street, city, state, and house number. You still do not need the who from the destination address. Previously, you ignored the issue of what to do if the source address was on a different street. Figure 1-9 shows how you can handle local off-street mail. Because it doesn't make sense to have a separate post office act as a switch for every street (it does not scale), you assign a number of streets in an area to a single post office. Each post office has access to more than one street, and acts as the switch when sending mail between houses on the same street.

Figure 1-9 Mail Switched Between Houses on the Same Street and Routed Between Houses on Different Streets



When a letter is sent from 1 Fullerton St. to 4 Fullerton St., the post office examines the destination street, city, and state names. The destination street, city, and state names are the same as the source street, city, and state names; so the letter is switched back to the proper house on Fullerton Street. If a letter is sent from 2 Diversey St. to 3 Kostner St., the local switch determines that it is not connected to Kostner Street. In this case, the switch forwards the letter to the FDK post office. Although the switch is located at the FDK post office, you can think of the post office having a separate switch and router department.

Upon receipt of the letter by the FDK post office, the source address is examined to make a decision on where to route the letter. Therefore, the FDK post office is acting as a router. Because the FDK post office has direct access to the street switches, the FDK post office is an **access router**. The layer in the mail delivery process containing the access router is called the **access layer**.

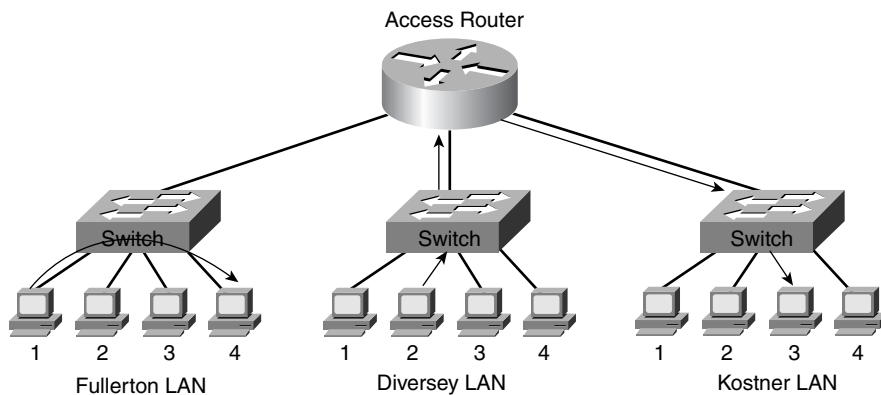
Just as the switches used a switching table, a router needs a routing table to determine how to route a particular piece of mail. The FDK post office is connected to three streets, and we will call the connection to these streets an interface. (See Table 1-2.)

Table 1-2 FDK Post Office Routing Table

Street	City	State	Interface
Fullerton	Chicago	Illinois	1
Diversey	Chicago	Illinois	2
Kostner	Chicago	Illinois	3

When a letter arrives at the FDK post office, the letter is routed on the street name and sent to the appropriate switch through the proper interface. The switch then switches the letter to the proper house using the house number. Off-street mail delivery is analogous to how multiple LANs can be connected using routers. If you replace the FDK post office with a router, you will have the configuration necessary for delivering messages between computers on different LANs in the same location. In Figure 1-10, an access router has replaced the FDK post office, and the computer network counterparts have replaced their respective switches.

Figure 1-10 Router Enables Communication Between LANs



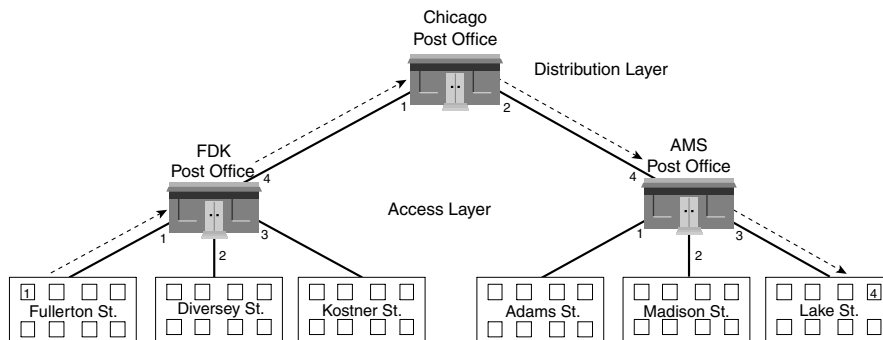
As with mail delivery, when computer 1 on the Fullerton LAN sends a message to computer 4 on the Fullerton LAN, the message is sent to the local switch. Because both computers are on the same LAN, the switch sends the message out the proper port to computer 4. When computer 2 on the Diversey LAN sends a message to computer 3 on the Kostner LAN, the switch forwards the message to the router because the destination is on a different LAN. The access router is directly connected to the switch that services the Kostner LAN, so the router forwards the message out the selected interface to the appropriate switch that switches the message out the proper port to computer 3.

Now you are able to use your off-street mail delivery scheme to develop a city-wide mail delivery system.

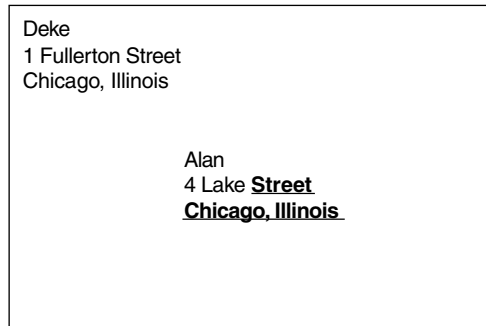
Citywide Delivery

A citywide mail delivery system can be realized by connecting the access post offices together using a *distribution post office*. (See Figure 1-11.)

Figure 1-11 System for Citywide Mail Delivery



To see how the distribution layer operates, you will trace the delivery of the letter shown in Figure 1-12.

Figure 1-12 Street, City, and State Are Used to Make a Delivery Decision

The letter is sent to FDK, the local access post office. Because the destination street is not connected to one of the FDK interfaces, the FDK post office must send the letter to the Chicago distribution post office. How does FDK know to send the letter to the Chicago distribution post office? There must be a rule telling FDK that if it receives a letter for a street that it is not connected to, by default, the letter is sent to the distribution post office. But what if the letter has a destination address that is not in the same city or state? The same rule applies. Any destination address that an access post office is not aware of is sent to the distribution post office. By default, any destination address that is not on Fullerton, Diversey, or Kostner streets is sent to the distribution post office.

For delivery to be successful between access post offices, the *default route* must be added to the routing tables for FDK and AMS. (See Tables 1-3 and 1-4.)

Table 1-3 FDK Post Office Routing Table

Street	City	State	Interface
Fullerton	Chicago	Illinois	1
Diversey	Chicago	Illinois	2
Kostner	Chicago	Illinois	3
Default	Default	Default	4

**Table 1-4** AMS Post Office Routing Table

Street	City	State	Interface
Adams	Chicago	Illinois	1
Madison	Chicago	Illinois	2
Lake	Chicago	Illinois	3
Default	Default	Default	4

When the Chicago post office receives the letter, a routing decision must be made based on the destination address. It is not sufficient to examine only the street name as the same street name can exist in different cities. Also, it is not sufficient to examine only the street and city names because there can be more than one city with the name Chicago. The distribution router must use the state, city, and street names to make a delivery decision—but not the street number or name of the recipient.

The Chicago post office knows about FDK and AMS because they are directly connected. To know about the street names that are serviced by FDK and AMS, Chicago needs a delivery or routing table similar to the routing table for the FDK and AMS post offices. (Refer to Tables 1-3 and 1-4.) The routing table for Chicago needs to contain a list of street names, and the interface to be used to reach these street names. (See Table 1-5.)

Table 1-5 Chicago Distribution Post Office Routing Table

Street	City	State	Interface
Fullerton	Chicago	Illinois	1
Diversey	Chicago	Illinois	1
Kostner	Chicago	Illinois	1
Adams	Chicago	Illinois	2
Madison	Chicago	Illinois	2
Lake	Chicago	Illinois	2

How do these values get into the routing table? The first way is to use static routes. A static route is manually entered into the routing table and does not change unless someone manually makes the change. The major problem with a static route is that there is no way to determine if the route actually exists. If the AMS post office burns down, Chicago thinks it still has a route to AMS because there is a static route in the routing table pointing to AMS. Mail could be delivered to a post office that is no longer there. Another method of placing routes in the routing table is to do so dynamically using a routing protocol. A routing protocol can be used to dynamically change the contents of the routing table when changes occur. We are not ready to discuss routing protocols just yet, so for now you will use only static routes.

Now consider the delivery of the letter with the Chicago routing table in place. FDK examines the street, city, and state names in the destination address. The street name is not known, so FDK consults its routing table and determines that the default says to use interface 4 and send the letter to the Chicago post office. The Chicago post office consults its routing table, looking for an entry that matches the street, city, and state names in the destination address. A match is made, and the letter is sent to the AMS post office using interface 2. AMS receives the letter, and routes it out interface 3 to the Lake Street switch. The Lake Street switch then switches the letter to the destination based on the street number.

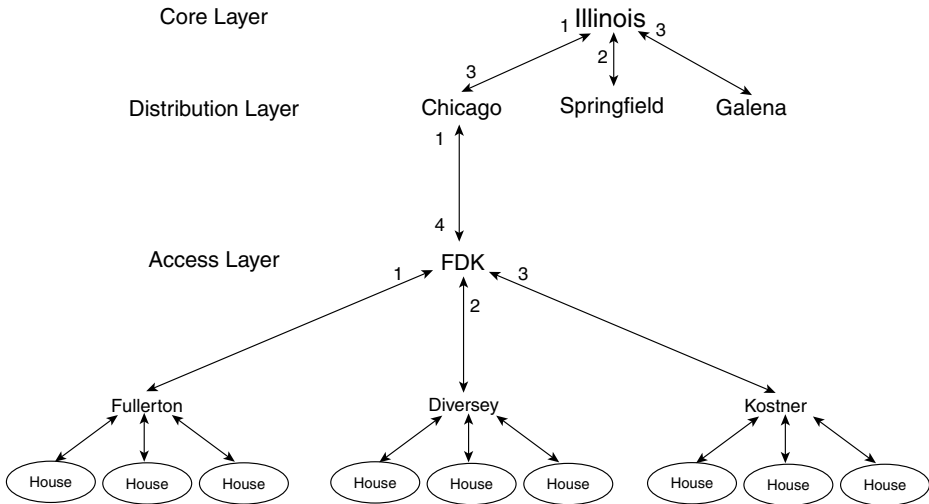
You created a citywide mail delivery system by adding another layer of delivery, or routing, called the distribution layer. Using the same process, you can now build a statewide mail delivery service by adding another routing layer to our system.

Statewide Delivery

Statewide delivery is for mail having the same state in both the destination and source addresses. The information required to deliver the mail is the city name, street name, and street number. Who receives the mail is still not used. To enable statewide delivery, another routing layer is added above the distribution layer. (See Figure 1-13.) This additional layer is called the core layer and is used for interstate and statewide delivery. For clarity, the AMS post office is not shown in Figure 1-13, although the routes for AMS are included in the routing table. (See Table 1-6.)



Figure 1-13 Core Delivery Layer Is Used for Statewide Delivery



When the *core post office* is added for the state of Illinois, a default route needs to be added to the routing table for the Chicago post office (See Table 1-6.)

Table 1-6 New Chicago Distribution Post Office Routing Table

Street	City	State	Interface
Fullerton	Chicago	Illinois	1
Diversey	Chicago	Illinois	1
Kostner	Chicago	Illinois	1
Adams	Chicago	Illinois	2
Madison	Chicago	Illinois	2
Lake	Chicago	Illinois	2
Default	Chicago	Illinois	3

When a letter is sent between cities in the same state, both the access and distribution routers will use their default routes and the letter will be delivered to the core router for the state. At this point, the function of the core router is to route mail between cities in the same state, so the only information that is required in the routing table is the city name (assuming the cities are in the state of Illinois).

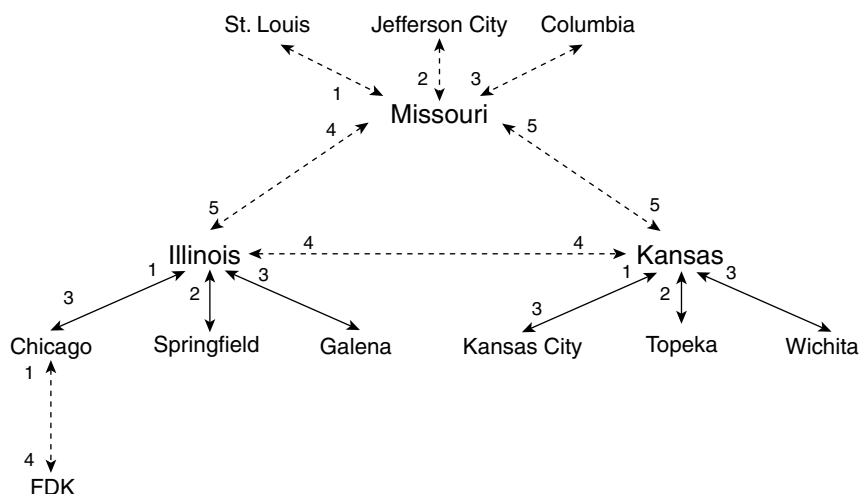
Table 1-7 State of Illinois Post Office Routing Table

City	State	Interface
Chicago	Illinois	1
Springfield	Illinois	2
Galena	Illinois	3

The core router for the state of Illinois will only examine the city and state information in the destination address to make a delivery or routing decision. Now that you have a scalable model for statewide delivery of the mail, you can easily expand your system to handle mail delivery for the country by connecting the core routers for each state.

Countrywide Delivery

Figure 1-14 shows a partial view of our countrywide mail delivery system. The core routers for Illinois, Missouri, and Kansas have been interconnected to enable the delivery of mail between these states.

Figure 1-14 Core Layer Is Used for Interstate Mail Delivery

Previously, when you added a new layer of routing, you had to add a default route to the routing table of the post office at the next lowest layer. For the core routers to deliver mail between states, you need to add specific routes for specific states (as shown in Table 1-8) for the state of Illinois post office. Now that there are two routes from Illinois to Kansas, and two routes from Illinois to Missouri (a direct connection to the state and an indirect connection through the other state), you need to determine how to handle multiple routes. The best connection to Kansas is directly to Kansas. But if that route is unavailable because of a bridge or road being out, you need to be able to use the other route as a backup. You can accomplish this by adding a cost to the route, and using the route with the lowest cost. If the route with the lowest cost is unavailable, use the other route. The cost of a route could represent distance, time, or any other factor that can be used to determine if one route is better than another route.

Table 1-8 State of Illinois Post Office Routing Table

State	City	Interface	Cost
Illinois	Chicago	1	1
Illinois	Springfield	2	1
Illinois	Galena	3	1
Kansas	Not Needed	4	1 (Best)
Kansas	Not Needed	5	2
Missouri	Not Needed	5	1 (Best)
Missouri	Not Needed	4	2

Notice in the routing table for the state of Illinois, as shown in Table 1-8, that less specific information is needed for delivery between states than is needed for delivery within the state. Mail destined for a different state than the core post office does not need to know the city name.

Finally, our plan for delivering mail throughout the United States is almost complete. Let's revisit how much information is required at each level of the system: core, distribution, and access.

At the beginning of the chapter, you decided that every post office did not need to know about every destination because this approach will not scale. Why is that?

If the core router for the state of Illinois needed to know about every destination, the routing table would be enormous. Every city in every state in the United States would have an entry in the routing table associating each city/state combination with an interface that is used to forward the mail along the proper route. If a new city is built or an old one ceases to exist, every routing table in every core post office would need to change, and this is not feasible. What you have achieved with your scheme is called **aggregation**, which is also known as **information hiding** or **summarization**.



Aggregation means that at the core post office level, the routing table entry for an interstate destination contains only the state information and an output interface. The core router does not need to maintain a route for every destination in the state. The core post office does need to maintain a route for cities in the state it services, and not the final destination in each of those cities. This makes the core routing table scalable; and it hides information regarding changes to destinations in other states, or connected cities. Therefore, if new houses are built or people move, the core post office does not care. What it needs to know is how to get mail to that state or city—and not who lives there or where they live.

Distribution post offices need only to know about the access-level post offices they are connected to and what streets they serve. As you have seen, a default route can be used to reach other destinations.

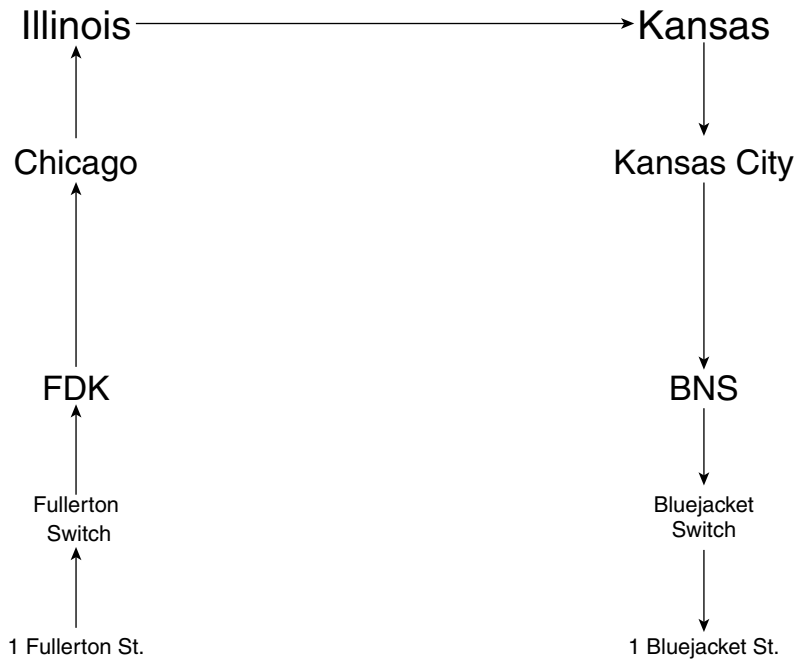
At the lowest level, the access routers maintain routes for streets that they service and have a default route for all other destinations.

Our system relies on static routing table entries to move the mail along in the proper direction. As the system grows, this will become an administrative nightmare. Eventually, you need to replace static routing with dynamic routing, but you have to wait a few chapters before we get there.

Before moving on to the next familiar system, you will trace a letter from Illinois to Kansas to see how the various layers of our postal delivery system operate. You will use a trace letter as shown in Figure 1-15, and see how it moves through our system illustrated in Figure 1-16.

Figure 1-15 Letter from Chicago to Kansas City

<p>Kankakee 1 Fullerton St. Chicago, Illinois</p> <p>Elvis <u>1 Bluejacket Street</u> <u>Kansas City, Kansas</u></p>
--

Figure 1-16 Tracing a Letter from Chicago to Kansas City

The following describes the process for delivering the letter from Illinois to Kansas:

1. Deliver the letter from Fullerton Street to the Fullerton switch.
2. The Fullerton switch must make a delivery decision on the street, city, and state names:

Elvis

1 Bluejacket Street

Kansas City, Kansas

3. The Fullerton switch does not know the street name, so it forwards to the FDK post office.