

CCIE职业发展系列
CCIE Professional Development

TCP/IP路由技术

(第一卷) (第二版) (英文版)

Routing TCP/IP
Volume I, Second Edition

A detailed examination of interior routing protocols

[美] **Jeff Doyle**, CCIE #1919
Jennifer Carroll, CCIE #1402 著

人民邮电出版社
北 京

内 容 提 要

本书是一本详细而又完整地介绍互连网络内部网关协议 (IGP) 的专业书籍, 堪称有关 IGP 方面不可多得的经典之作。本书共分三个部分。第一部分主要介绍了网络和路由选择的基本知识, 其中包括 IPv4 协议、IPv6 协议和路由技术。第二部分是本书的精华, 这一部分详细、深入地讲述了各种常用的内部路由协议, 如 RIP、RIPv2、RIPng、无类别路由选择、EIGRP、OSPFv2、OSPFv3、IS-IS 等协议, 每一章除了对该协议的实现机制和参数详尽阐述, 使读者对协议的实现原理有一个清晰的理解外, 还通过在实际网络环境中的实例, 详细地论述了该协议在 Cisco 路由器上的配置和故障处理方法, 帮助读者获取大量解决实际问题的专业技能。第三部分介绍了如路由重新分配、缺省路由/按需路由选择、路由过滤、路由映射等多种重要而有效的路由控制工具, 用来创建和管理多个 IP 路由选择协议的协调和互操作。附录部分讲述了二进制、十六进制转换、访问列表、CCIE 提示等内容。

相对于第一版, 本书第二版具有以下更新: 在第一版详细讲述 IPv4 协议中 IGP 的基础上, 大量增加了相应协议在 IPv6 协议中的实现和配置, 其中单独一章用来讲述 IPv6 中应用的 OSPFv3 协议, 这是本书新版的一大亮点; 同时本书根据 Internet 和 Cisco IOS 系统的最新发展, 适当地删减了如网桥、IGRP 等过时的内容, 并增加了许多新的 IOS 增强特性的讲解。

本书的读者不仅是那些准备 CCIE 考试的考生, 也是任何需要完整理解 IPv4, 特别是 IPv6 下 TCP/IP 内部路由选择协议实现的网络设计和工程人员。本书中对协议细节的讲解和对网络实例的探讨相信会让读者获益匪浅。

About the Authors

Jeff Doyle (CCIE No. 1919) specializes in IP routing protocols, MPLS, and IPv6. He has designed or assisted in the design of large-scale IP service provider networks throughout North America, Europe, Japan, Korea, and the People's Republic of China. Jeff has presented numerous corporate seminars, and has also spoken at NANOG, JANOG, APRICOT, and at IPv6 Forum conferences. Jeff holds a BA from Memphis State University, and studied Electrical Engineering at the University of New Mexico. Jeff lives in Denver, Colorado.

Jennifer Carroll (CCIE No. 1402) is an independent network consultant in Redmond, Washington. She has designed, implemented, and optimized many TCP/IP networks, and has developed and taught a variety of networking and internetworking courses on routing protocols and Cisco routers over the past 15 years. Jennifer can be contacted at jennifer.carroll@ieee.org.

About the Technical Reviewers

Frank Knox, Chief Technical Officer, has been with Skyline Computer for a little over six years. He is a dual CCIE (CCIE No. 3698: SNA/IP and Routing/Switching) as well as a CCSI. In addition to his CTO responsibilities, Frank teaches several advanced Cisco-related courses, including a one-week CCIE Lab Preparation Workshop. He is considered to be an expert in mainframe attached router technologies and in the technologies and issues associated with integrated networking (for example, SNA/IP and Voice/Data). He has more than 37 years of networking experience with IBM, GTE (Verizon) Directories, and Skyline Computer Corp. This experience includes field service, field support, product planning, management, and all facets of networking education. In addition, he developed and taught several courses for the University of Dallas Telecommunications MBA program. Frank also has an MS degree in Telecommunications from Pace University (4.0 GPA).

After working in various roles as an engineer within Cisco for the past 6.5 years, **Steven Edward Moore** transitioned to the IP Routing Protocol Scalability Team. There, his focus encompasses all aspects of extending network and protocol scalability: considering new features and optimizations to the protocol architecture, designing tests to measure current protocol scalability, working with customers to implement scaling functionality in their network, and participating in events such as Networkers to educate others on how to enhance their network's performance and scalability from the routing perspective.

Rena Yang is a software engineer at Cisco Systems. She has more than six years of experience implementing code in Cisco IOS. She currently works on IS-IS. Before this, she focused on IPv4, UDP, access lists, policy routing, and routing infrastructure. Rena holds a bachelor's of science and masters of engineering in computer science from MIT.

Dedications

I would like to dedicate this book to my wife, Sara, and my children, Anna, Carol, James, and Katherine.

—Jeff

I would like to dedicate this book to my husband, Mike, and sons, Mitchell and Jonathan. Their patience and support helped me complete this book.

—Jennifer

Acknowledgments

Many thanks to Brett Bartow, Chris Cleveland, Andrew Cupp, San Dee Phillips, and all of the staff of Cisco Press who made this book possible.

The technical editors, Steven Moore, Rena Yang and Frank Knox, did a fantastic job. We want to thank them for their outstanding advice and recommendations.

We want to thank Frank Knox, Carl Pike, Chris Tonini, and the rest of the employees of Skylabs networks. Skylabs' lab setup and access to the lab is easy to use and had everything we needed to complete all the configurations and case studies in this book.

Foreword

In 1976, when I saw my first Arpanet IMP at Digital Equipment Corporation, networks as we know them today were in their infancy. SNA, XNS, and DECnet were under early development, and packet switching versus circuit switching was the hot topic of the day. Those of us involved in the design of the switching and routing algorithms were dealing with routers (although we didn't call them that) that had 64 kilobytes of memory, data link of 56 kilobits were considered blindingly fast, and networks with 256 nodes were big enough that if you were the salesman who sold those 256 computers, you would retire fabulously wealthy.

Thirty years is a long time, and today the individual networks that make up the Internet contain thousands or tens of thousands of nodes, while the Internet as a whole contains hundreds of millions of computers. Most striking in the evolution over this human generation is that the foundations of the Internet laid down in the TCP/IP protocol suite have survived mostly intact through four or more generations of computing architectures, three complete generations of operating system technology, and an increase of five orders of magnitude in transmission speeds.

Yet, we still treat routing in packet-switched networks as a black art. Why is that?

First, designing robust, scalable distributed algorithms is hard. Despite our best intentions to make them simple, complexity creeps in to deal with the inevitable special cases, optimizations, peculiar topologies, and link technologies one encounters. Because a "fork lift upgrade" of an entire network is rarely feasible, we have multiple generations of technology present simultaneously, and we must maintain backward-compatibility with essentially no disruption to deployed services. As policies governing the routing of packets become more sophisticated, our ability to devise automated discovery and configuration procedures gets overwhelmed, and we fall back on manual configuration and performance tuning techniques. Finally, as the environment in which these networks are operated has evolved from a cooperative one where trust was implicit to one in which the network is subject to both inside and outside attack, designing and deploying routing systems that can be made secure has become an urgent priority.

Routing TCP/IP tackles this black art comprehensively. The present Volume 1 covers all the needed fundamentals of TCP/IP networks and gives you all the tools needed to understand how routing is accomplished within a single administrative region of the Internet. Straightforward ideas of packet-switched routing are presented first in the chapters on addressing and static routing. The most popular IGP's—RIP, EGRP, OSPF, and ISIS—are covered in depth. Advanced topics in route redistribution, route filtering, and policy routing round out Volume 1.

This second edition also adds essential material on IPv6 as well as bringing all the material up to date with examples and configurations for the latest releases of Cisco IOS.

For anyone wanting a comprehensive understanding of how routing in TCP/IP networks really works, from the design principles of routing algorithms, to the evolution of addressing schemes, to the practical aspects of designing and configuring the routing of large autonomous systems, this is the book for you.

David Oran
Cisco Fellow

Introduction

Routing is an essential element of all but the smallest data communications networks. At one level, routing and the configuration of routers are quite simple. But as networks grow in size and complexity, routing issues can become at once both large and subtle. Perversely, perhaps, we are grateful for the difficult problems large-scale routing can present—as network systems consultants, these problems are our bread and butter. Without them, the phrase “You want fries with that?” could be an unfortunate part of our daily vocabulary.

Cisco Certified Internetwork Experts are widely recognized for their ability to design, troubleshoot, and manage large networks. This recognition comes from the fact that you cannot become a CCIE by attending a few classes and then regurgitating some memorized facts onto a written test. A CCIE has proven expertise in an intense, famously difficult hands-on lab exam.

Objectives

This book is the first of two volumes that focuses on TCP/IP routing issues. Early in the writing of the first edition, Kim Lew, former Cisco Systems program manager, said, “Our objective is to make CCIEs, not to make people who can pass the CCIE lab.” We entirely agree with that statement and have used it as a guiding principle throughout the writing of this book. Although the book includes many case studies and exercises to help you prepare for the CCIE lab, my primary objective is to increase your understanding of IP routing—both on a generic level and as it is implemented on Cisco routers.

Audience

The audience for this book is any network designer, administrator, or engineer who needs a full understanding of the interior routing protocols of TCP/IP. Although the practical aspects of the book focus on the Cisco IOS, the information is applicable to any routing platform.

The book is not only for readers who plan to become CCIEs, but for people who wish to advance their knowledge of TCP/IP routing. These readers will fall into one of three categories:

- The “beginners” who have some basic networking knowledge and wish to begin a deep study of networking.
- The intermediate-level networking professionals who have experience with routers, Cisco or otherwise, and plan to advance that experience to the expert level.
- The highly experienced networking experts. These individuals have extensive hands-on expertise with Cisco routers and are ready to take the CCIE lab; however, they want a structured review and series of exercises for verification and validation.

CCIE Professional Development: Routing TCP/IP, Volume I focuses primarily on intermediate-level networking professionals while offering to beginners a structured outline of fundamental information and to experts the required challenges to hone their skills.

Changes from First Edition

There are several factors influencing the changes contained in this second edition. The first factor is the CCIE itself. When I (Jeff) wrote the first edition of this book, the CCIE—specifically what is now

called the Routing and Switching specialty of the CCIE—was the only certification Cisco Systems offered. Now, there is a series of certifications creating a path to the CCIE at the pinnacle. Moreover, the typical networking professional is more knowledgeable than in 1997. Given this, we have eliminated the first chapter of the original book, which covered such very basic concepts as the definition of bridges and routers and network addresses. (When was the last time you even saw a bridge in a network?)

The second factor influencing the changes in this edition is the changes in the Cisco Systems IOS. IGRP, which was frequently used when the first edition was written, is now a legacy protocol whose main significance is as the ancestor of EIGRP. Therefore the IGRP chapter of the first edition has been eliminated and IGRP is covered for historical perspective early in the EIGRP chapter. The IOS command suite itself has expanded to accommodate new functions and options; we have made every effort to include the commands and protocol extensions that did not exist in the late 1990s.

Lastly, a protocol that existed mostly only in proposal form in 1997—IPv6—is now in the early stages of worldwide deployment. You can expect to need a detailed knowledge of this protocol and the extensions to IP routing protocols that support it in the near future, if not already, so this second edition delves deeply into routing IPv6.

Other changes in this edition are semantic. For example, in the first edition, I (Jeff) made a point of differentiating between a “network” as a data link and an “internetwork” as a set of networks connected by routers. Although that terminology is certainly accurate, it is clumsy, and “internetwork” is seldom used these days. Instead, “network” usually refers to everything from a local link to worldwide autonomous systems operated by the likes of Level 3, NTT, and Sprint. We have attempted to bring the terminology in this edition up to modern, common usage.

Organization

The 14 chapters of the book are divided into three parts.

Part I, “Routing Basics,” examines the basics of IPv4 and IPv6, and the basics of routing. Although more advanced readers may wish to skip the first chapter, we recommend that they at least skim Chapter 3, “Static Routing,” and Chapter 4, “Dynamic Routing Protocols.” And, of course, if you are not yet familiar with IPv6, Chapter 2, “IPv6 Overview,” is a must-read.

Part II, “Interior Routing Protocols,” covers the IP Interior Gateway Protocols. Each protocol-specific chapter begins with a discussion of the theory, mechanics, and parameters of the protocol. This general overview is followed by case studies on configuring and troubleshooting the protocol using Cisco Systems’ IOS in various network topologies.

The Exterior Gateway Protocol, BGP, and topics such as multicast routing, Quality of Service, router security and management, and Network Address Translation, are covered in “Routing TCP/IP, Volume II.”

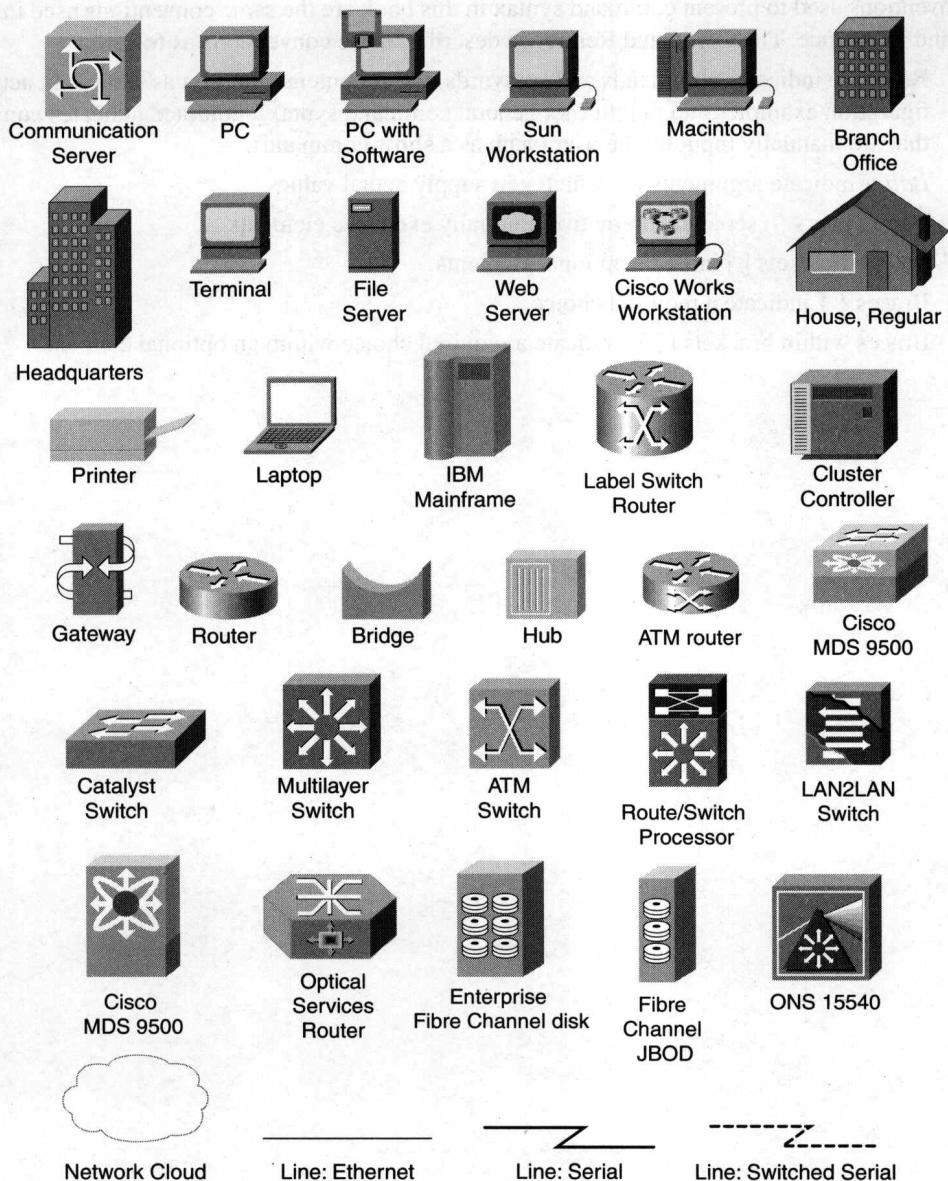
Part III, “Route Control and Interoperability,” examines the tools available for creating and managing interoperability with multiple IP routing protocols, and also such tools as default routes and route filtering. As such, the chapters of this last part provide an introduction to the tools necessary for building the complex routing policies introduced in Volume II. These chapters, like the ones in Part II, begin with concepts and conclude with case studies.

Book Features

Most chapters conclude with a set of review questions, configuration exercises, and troubleshooting exercises. The review questions focus on the theoretical aspects of the chapter topic, whereas the configuration and troubleshooting exercises address Cisco-specific aspects of the chapter topic.

Also at the end of each chapter is a table with a brief description of all important Cisco IOS commands used in that chapter. The conventions used to present these commands are the same conventions used in the IOS Command Reference and presented earlier in this introduction.

Icons Used in This Book



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italics* indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

Contents at a Glance

Part I	Routing Basics	3
Chapter 1	TCP/IP Review	5
Chapter 2	IPv6 Overview	49
Chapter 3	Static Routing	81
Chapter 4	Dynamic Routing Protocols	131
Part II	Interior Routing Protocols	167
Chapter 5	Routing Information Protocol (RIP)	169
Chapter 6	RIPv2, RIPv6, and Classless Routing	201
Chapter 7	Enhanced Interior Gateway Routing Protocol (EIGRP)	251
Chapter 8	OSPFv2	335
Chapter 9	OSPFv3	477
Chapter 10	Integrated IS-IS	513
Part III	Route Control and Interoperability	629
Chapter 11	Route Redistribution	631
Chapter 12	Default Routes and On-Demand Routing	677
Chapter 13	Route Filtering	701
Chapter 14	Route Maps	731
Part IV	Appendixes	767
Appendix A	Tutorial: Working with Binary and Hex	769
Appendix B	Tutorial: Access Lists	775
Appendix C	CCIE Preparation Tips	803
Appendix D	Answers to Review Questions	809
Appendix E	Solutions to Configuration Exercises	827
Appendix F	Solutions to Troubleshooting Exercises	881
Index		887

Table of Contents

Part I Routing Basics 3

Chapter 1 TCP/IP Review 5

TCP/IP Protocol Layers	5
IP Packet Header	7
IPv4 Addresses	16
First Octet Rule	18
Address Masks	20
Subnets and Subnet Masks	22
Designing Subnets	24
Breaking the Octet Boundary	25
Troubleshooting a Subnet Mask	29
Address Resolution Protocol (ARP)	30
Proxy ARP	35
Gratuitous ARP	36
Reverse ARP	36
Internet Control Message Protocol (ICMP)	37
Host-to-Host Layer	41
TCP	41
UDP	44
Looking Ahead	45
Summary Table: Chapter 1 Command Review	45
Recommended Reading	46
Review Questions	46
Configuration Exercises	47
Troubleshooting Exercises	47

Chapter 2 IPv6 Overview 49

IPv6 Addresses	50
Address Representation	50
IPv6 Address Types	52
Global Unicast Addresses	53

Identifying IPv6 Address Types	54
Local Unicast Addresses	55
Anycast Addresses	55
Multicast Addresses	56
Embedded IPv4 Addresses	58
IPv6 Packet Header Format	58
Extension Headers	60
ICMPv6	63
Neighbor Discovery Protocol	64
NDP Messages	65
Router Discovery	70
Address Autoconfiguration	71
Duplicate Address Detection	73
Neighbor Address Resolution	74
Privacy Addresses	75
Neighbor Unreachability Detection	76
Looking Ahead	77
Review Questions	77
Chapter 3 Static Routing	81
Route Table	81
Configuring Static Routes	84
Case Study: Simple IPv4 Static Routes	85
Case Study: Simple IPv6 Static Routes	90
Case Study: Summary Routes	94
Case Study: Alternative Routes	96
Case Study: Floating Static Routes	97
Case Study: IPv6 Floating Static Routes	100
Case Study: Load Sharing	102
Load Sharing and Cisco Express Forwarding	103
Per Destination Load Sharing and Fast Switching	105
Per Packet Load Sharing and Process Switching	105
Which Switching Method Will Be Used?	106
Case Study: Recursive Table Lookups	107
Troubleshooting Static Routes	109
Case Study: Tracing a Failed Route	109
Case Study: A Protocol Conflict	112
Case Study: A Replaced Router	116
Case Study: Tracing An IPv6 Failed Route	120
Looking Ahead	122

Summary Table: Chapter 3 Command Review	123
Review Questions	124
Configuration Exercises	124
Troubleshooting Exercises	126
Chapter 4 Dynamic Routing Protocols	131
Routing Protocol Basics	131
Path Determination	132
Metrics	133
Hop Count	134
Bandwidth	134
Load	135
Delay	135
Reliability	135
Cost	135
Convergence	135
Load Balancing	137
Distance Vector Routing Protocols	137
Common Characteristics	138
Periodic Updates	138
Neighbors	138
Broadcast Updates	138
Full Routing Table Updates	139
Routing by Rumor	139
Route Invalidation Timers	140
Split Horizon	141
Counting to Infinity	143
Triggered Updates	144
Holddown Timers	144
Asynchronous Updates	145
Link State Routing Protocols	146
Neighbors	147
Link State Flooding	147
Sequence Numbers	148
Aging	153
Link State Database	154
SPF Algorithm	156
Areas	160
Interior and Exterior Gateway Protocols	160
Static or Dynamic Routing?	162

Looking Ahead	164
Recommended Reading	164
Review Questions	164

Part II Interior Routing Protocols 167

Chapter 5 Routing Information Protocol (RIP) 169

Operation of RIP	170
RIP Timers and Stability Features	170
RIP Message Format	172
Request Message Types	174
Classful Routing	175
Classful Routing: Directly Connected Subnets	176
Classful Routing: Summarization at Boundary Routers	176
Classful Routing: Summary	178
Configuring RIP	178
Case Study: A Basic RIP Configuration	178
Case Study: Passive Interfaces	181
Case Study: Configuring Unicast Updates	182
Case Study: Discontiguous Subnets	184
Case Study: Manipulating RIP Metrics	187
Case Study: Minimizing the Impact of Updates	190
Troubleshooting RIP	193
Looking Ahead	193
Summary Table: Chapter 5 Command Review	193
Recommended Reading	194
Review Questions	194
Configuration Exercises	194
Troubleshooting Exercises	196

Chapter 6 RIPv2, RIPv6, and Classless Routing 201

Operation of RIPv2	201
RIPv2 Message Format	202
Compatibility with RIPv1	205
Classless Route Lookups	205
Classless Routing Protocols	206
Variable-Length Subnet Masking	206
Authentication	210

Operation of RIPng	212
Configuring RIPv2	214
Case Study: A Basic RIPv2 Configuration	215
Case Study: Compatibility with RIPv1	215
Case Study: Using VLSM	218
Case Study: Discontiguous Subnets and Classless Routing	220
Case Study: Authentication	222
Configuring RIPng	224
Case Study: Basic RIPng Configuration	224
Case Study: RIPng Process Customization	232
Case Study: Metric Manipulation	233
Case Study: Route Summarization	235
Troubleshooting RIPv2 and RIPng	236
Case Study: Misconfigured VLSM	236
Looking Ahead	242
Summary Table: Chapter 6 Command Review	242
Recommended Reading	244
Review Questions	244
Configuration Exercises	244
Troubleshooting Exercises	246

Chapter 7 Enhanced Interior Gateway Routing Protocol (EIGRP) 251

The Roots of EIGRP: An Overview of IGRP	251
Process Domains	252
IGRP Timers and Stability Features	254
IGRP Metrics	255
From IGRP to EIGRP	261
Operation of EIGRP	263
Protocol-Dependent Modules	263
Reliable Transport Protocol	264
Neighbor Discovery/Recovery	265
Diffusing Update Algorithm	266
DUAL: Preliminary Concepts	267
DUAL Finite State Machine	273
Diffusing Computation: Example 1	276
Diffusing Computation: Example 2	280