

CISCO系列丛书（影印版）



Robert Wright

A thorough guide to understanding IP
behavior in a router-based network

IP ROUTING PRIMER

IP路由技术基础

CISCO SYSTEMS

CISCO PRESS

清华大学出版社

<http://www.tup.tsinghua.edu.cn>



IP ROUTING PRIMER

IP 路由技术基础

Robert Wright

Robert Wright



清华大学出版社

(京)新登字 158 号

IP Routing Primer

Robert Wright

“Authorized reprint from the English language edition published by Macmillan Technical Publishing
Copyright © 1998”

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 retrieval system, without permission from the Publisher.

English language reprint edition published by Tsinghua University Press
Copyright © 1999”

This edition is authorized for sale only in the following Territory: The People's Republic of China, with the exception of the Hong Kong and Macau territories.

本书封面贴有清华大学出版社激光防伪标签,无标签者不得销售。

北京市版权局著作权合同登记号: 01-99-0592

图书在版编目(CIP)数据

IP 路由技术基础: 英文/(美)赖特(Wright, R.)著. 影印版. - 北京: 清华大学出版社, 1999.4

(CISCO 系列丛书)

ISBN 7-302-03456-7

I. I… II. 赖… III. 因特网-通信协议-英文 IV. TP393.4

中国版本图书馆 CIP 数据核字(1999)第 08649 号

出版者: 清华大学出版社(北京清华大学校内, 邮编 100084)

<http://www.tup.tsinghua.edu.cn>

印刷者: 清华大学印刷厂

发行者: 新华书店总店北京发行所

开 本: 787×960 1/16 印张: 20

版 次: 1999 年 4 月第 1 版 1999 年 6 月第 2 次印刷

书 号: ISBN 7-302-03456-7/TP·1883

印 数: 3001~7000

定 价: 30.00 元

出版前言

21 世纪将会是一个信息高速公路四通八达的时代,信息产业发展的水平亦将是评估一个国家综合国力的重要依据。世界各国将会在信息技术的研究开发和信息产业的发展方面展开激烈的竞争。这既是一种挑战也是一种机遇。有鉴于此,我国已经开始全面快速地发展网络技术和因特网。

Cisco Systems 公司是世界领先的全球 Internet, Intranet 以及电信网络设备及解决方案的供应商,1996 年名列世界十大电信公司之一。Cisco Systems 的联网操作系统(IOS)是支持网络服务和网络应用的坚实基础。该公司与麦克米伦计算机出版公司合作创立了 Cisco Press,出版了一系列关于最新的网络技术的权威著作。这些著作不仅兼顾建网与网际互联的基础理论和实际应用,为网络专业人员和用户提供必要的技术支持,还有一部分是为 Cisco CCIE 考试和 CCNA,CCNP,CCDA 及 CCDP 职业考试认证准备的自学和培训教材。Cisco 公司早于 1994 年就进入中国,已为国内信息产业界所熟悉。我们引进其中部分著作组成“CISCO 系列丛书(影印版)”影印出版,以祈对我国信息产业的发展稍尽绵薄之力,并衷心希望这套丛书对从事建网,网际互联的专业人员;有志于我国信息产业发展的读者,以及参加 Cisco 培训和准备 Cisco 考试认证的人员有所裨益。

清华大学出版社

1999.4



Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Macmillan Technical Publishing 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.



About the Author

Robert Wright is Manager of Customer Engineering for an Internet backbone startup company. He was formerly a Senior Manager with Cisco Systems, Inc. He managed a team of Senior Customer Engineers responsible for providing dedicated, highly individualized technical support for some of the world's largest and fastest growing Internet service providers. He has also worked in Cisco's Critical Accounts group, in Cisco's Technical Assistance Center (TAC), and as a Systems Engineer for Cisco in the United Kingdom. Prior to joining Cisco, Robert worked for Texaco Ltd. in the United Kingdom. He worked in Texaco's LAN support group and supported the Cisco routers in the UK. He was also part of the team responsible for designing and installing Texaco's new data center in East London. He currently resides in San Jose, California, and enjoys traveling with his wife, golfing, water skiing, snowboarding, and listening to his collection of Jimmy Buffett CDs.

A decorative line starts from the top left corner, extends diagonally down and to the right, then turns 90 degrees and extends horizontally to the right edge of the page.

Dedications

This book is dedicated to my wife, Lelia. She is always there for me and is my best friend.



Acknowledgments

I wish to acknowledge Tony Li and Robert Craig for their infinite patience and desire to spread their knowledge when I was first learning many of the concepts I have presented in this book. I would also like to thank the members of the Cisco ISP Expert Team who helped me validate many of the concepts presented in this material while I was developing it: Nga Vu for her perseverance in getting me to create the training document that was the genesis of this book and, finally, Shankar Vemulapalli and Srinivas Vegesna for their excellent technical editing. Finally, I would like to thank my wife, Lelia, for her support and motivation during the past year.



Introduction

Welcome to Cisco's *IP Routing Primer*! This book covers the generic behavior of IP routing and packet forwarding using Cisco routers. It goes into detailed analysis of several real-life scenarios to provide insight into the fundamentals of IP that everybody supporting IP in a network should know.

By providing examples taken directly from Cisco routers, this book enables the reader to associate theoretical behaviors discussed in many internetworking books with their real-life counterparts. The reader should find it much easier to understand statements such as "Split horizon refers to the concept of not advertising routes over the interface they were learned from . . ." when it is accompanied by actual screen output from a Cisco router as it sends a routing table update to an adjacent router.

By providing examples of IP routing behaviors taken directly from Cisco routers, this book avoids theoretical explanations that can vary from one networking engineer to another. Ask any 10 network engineers what a poison reverse update is, and you will likely receive 10 different answers. Refer to the section on Poison Reverse and Triggered Updates in Chapter 2 for a real-life example you can apply to any situation.

Objective of This Book

This book was written to fill the gap between internetworking books that are long on theory but short on practice and the high-level seminars on internetworking that cover real-life examples but lack depth. After reading this book, the reader should have a better understanding of the complexities involved in designing and supporting IP networks.

By presenting a few of the most common issues that can be encountered in managing an IP network, I hope to pique your curiosity to learn more by going into a networking lab and experimenting on your own.

Another goal of this book is to convey the idea that to be a successful internetworking engineer, it is more important to understand what needs to happen in a given situation to make something work than to memorize exactly what happens. For example, you can memorize the fact that before an IP host can send its first IP packet to another host, it must send an ARP request and receive an ARP reply. But what have you really learned? If you don't understand that the IP host was resolving a layer three address to a layer two address to avoid broadcasting all layer two frames to all hosts on the network, you haven't grasped a fundamental part of what makes a network work.

In addition, this book teaches some fundamental skills that anyone involved in internetworking should possess, such as binary and hexadecimal numbering and IP subnetting. It is almost impossible to be a successful network engineer if you do not have these skills mastered.

Keep in mind that this book is not intended to be an in-depth analysis of the individual routing protocols presented. This

information is readily available in many other books, RFCs, and white papers. The goal is to present the behaviors exhibited by a couple of routing protocols (RIP and IGRP) to set the stage for reading and understanding the material available from other sources.

In addition, this book is not intended to be a design and implementation guide for building IP networks. Instead, it is meant to be a guide to issues that must be addressed when building IP networks. It gives readers a better idea of which questions to ask and which problems need to be solved when building their own networks.

Audience

This book is intended for anybody involved in supporting or designing IP networks—engineers, support personnel, and the like. It covers many basic internetworking concepts that people just starting out need to understand. It goes into great detail on some very sophisticated topics that even those with several years of experience supporting IP networks will find interesting.

Conventions Used in This Book

The routers in the topology and configuration section may appear in different topologies throughout this book. With the exception of removing some links in certain examples, the actual configurations do not change from those shown unless specifically noted.

Router commands referenced in paragraphs are in lowercase and italics, for example, *show ip route*.

Screen output from routers is presented in a monospaced font. For example:

```
RouterA#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

    168.71.0.0/16 is subnetted, 5 subnets
C       168.71.9.0 is directly connected, Serial1
R       168.71.8.0 [120/1] via 168.71.9.2, 00:00:39, Serial1
R       168.71.7.0 [120/1] via 168.71.6.2, 00:00:11, Serial0
           [120/1] via 168.71.9.2, 00:00:39, Serial1
C       168.71.6.0 is directly connected, Serial0
C       168.71.5.0 is directly connected, Ethernet0
S*    0.0.0.0/0 is directly connected, Ethernet0
RouterA#
```

Occasionally, a specific portion of the screen output is referenced by the surrounding text. In these cases, the output in question will be in bold.


For example, note that the update timer for 168.71.8.0 in the following output of the *show ip route* command from RouterA is now 39 seconds.

```
RouterA#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route


Gateway of last resort is 0.0.0.0 to network 0.0.0.0

    168.71.0.0/16 is subnetted, 5 subnets
C       168.71.9.0 is directly connected, Serial1
R       168.71.8.0 [120/1] via 168.71.9.2, 00:00:39, Serial1
R       168.71.7.0 [120/1] via 168.71.6.2, 00:00:11, Serial0
           [120/1] via 168.71.9.2, 00:00:39, Serial1
C       168.71.6.0 is directly connected, Serial0
C       168.71.5.0 is directly connected, Ethernet0
S*    0.0.0.0/0 is directly connected, Ethernet0
RouterA#
```

In addition, this book utilizes two other conventions:



Margin notes—These are used to add additional points of interest to the reader without disrupting the flow of the main ideas presented.



Margin hints—These can be troubleshooting tips or additional procedures that are relevant to the current subject matter.

Organization

This book is divided into eight chapters and one appendix, as follows:

- Chapter 1, “Topology and Router Configurations”
This chapter presents the routers and their configurations, which will form the basis of the scenarios presented in this book. It also introduces some of the basic functions of a router and some of the problems a router must solve to do its job successfully.
- Chapter 2, “Routing Metrics and Distances”
This chapter explains what routing metrics are and how they can be calculated. It also describes Cisco’s utilization of the distance function to determine which routing protocols take precedence when they run concurrently.
- Chapter 3, “Discontiguous Networks, Summarization, and Subnet 0”
This chapter explains what the terms discontiguous networks, summarization, and subnet 0 mean and how

they interact in a live network. It also includes scenarios in which functions have been used incorrectly to show the problems they can cause.

- Chapter 4, “Using IP Unnumbered and VLSM”

This chapter describes IP unnumbered and Variable Length Subnet Masking (VLSM) and explains how they can be used as tools when building IP networks. It also includes scenarios in which these functions have been used incorrectly to show the problems they can cause.

- Chapter 5, “Default Routing”

This chapter explains what default routing is and why it is necessary. Several scenarios are provided to show how default routing works and what can happen when it is not configured properly.

- Chapter 6, “IP Troubleshooting Scenarios”

This chapter walks through common IP connectivity problems and introduces some tools and techniques to resolve them.

- Chapter 7, “Bridging IP Between Dissimilar Media”

Many network engineers have made the mistake of attempting to bridge IP between Token Ring and Ethernet using Cisco routers. Cisco routers do not support this function. This chapter explains why this is the case.

- Chapter 8, “Hexadecimal and Binary Numbering and IP Addressing”

This chapter describes the two numbering systems and explains why it is important to have mastered using them. It also covers IP addressing and subnetting.

- Appendix A, “RFCs”

This chapter includes all of the RFCs referenced in this book, as well as a few that are useful for people just getting started in internetworking. In addition, there are several references to RFCs on more advanced topics.

It is recommended that you start with Chapter 1 because the concepts build on one another as the book progresses. Welcome again to Cisco’s *IP Routing Primer*!

Contents

	Introduction	xxiii
Chapter 1	Topology and Router Configurations	1
	Understanding the Role of Routers in Networks	1
	The Router Interface	2
	Network Layer Addresses	2
	<i>Datagrams</i>	3
	<i>MAC Addresses</i>	4
	<i>IP Address Formats</i>	5
	Network Reference Models	6
	Understanding Topology and Router Configurations	9
	RouterA's Configuration	10
	RouterB's Configuration	11
	RouterC's Configuration	12
	Understanding What a Router Does	13
	Sample Network	13
	How a Router Knows What to Do	15
	Choosing Your Routing Protocol	16

Understanding How Forwarding Decisions Are Made	18
Performing Longest Match Lookups	18
Forwarding Decisions for Multipoint Interfaces	21
End Systems Sending Packets to Other Subnets	23
Summary	25

Chapter 2	Routing Metrics and Distances	27
	Primary Activities of Convergence	27
	Viewing the Invalid Timers in a Routing Table	29
	Viewing an Expired Invalid Timer in a Routing Table	30
	Router Still Uses a Path	33
	Understanding Convergence	37
	Parallel Paths	38
	<i>The Effect of Parallel Paths on Convergence</i>	38
	<i>Looking at Parallel Paths in a Routing Table</i>	39
	Convergence in Action	40
	<i>The Routing Table After Convergence</i>	41
	<i>Step-by-Step Review of Convergence</i>	43
	Debug Messages and Reality	46
	When Holddown Is Initiated	46
	Understanding Parallel Paths and Their Effect on Packet Forwarding	47
	Process Switching Versus Fast Switching	50
	Configuring Process Switching	51
	Configuring Fast Switching	52
	Understanding the Role of Split Horizon	54
	Routing Advertisements with Split Horizon Enabled	56
	Routing Advertisements with Split Horizon Disabled	57