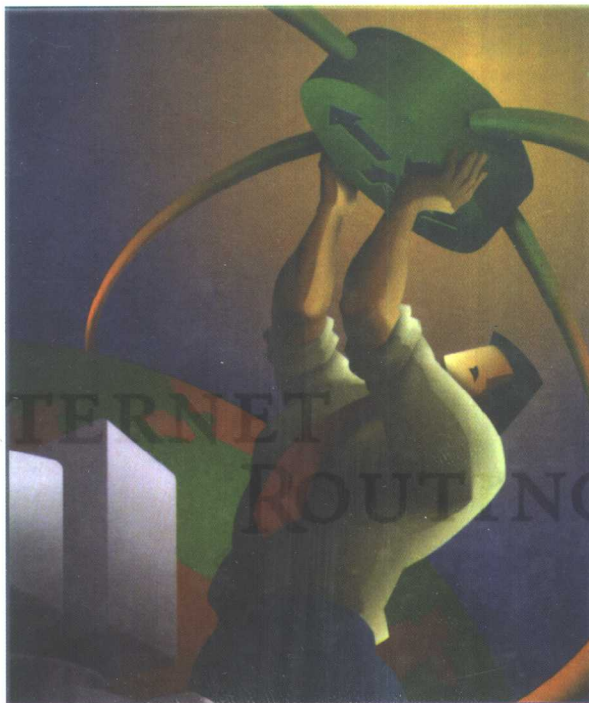


**CISCO**系列丛书（影印版）



Bassam  
Halabi

The definitive resource for internetworking  
design alternatives and solutions

# INTERNET ROUTING ARCHITECTURES

## INTERNET路由结构

CISCO SYSTEMS  
  
CISCO PRESS

清华大学出版社

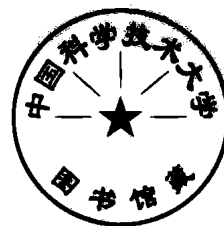
<http://www.tup.tsinghua.edu.cn>



**Internet Routing Architectures**

# **Internet 路由结构**

**Bassam Halabi**



**清华大学出版社**

(京)新登字 158 号

## Internet Routing Architectures

Bassam Halabi

“Authorized reprint from the English language edition published by New Riders Publishing  
Copyright © 1997”

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: 01-2000-0014

书 名: Internet 路由结构

作 者: Bassam Halabi

出版者: 清华大学出版社(北京清华大学学研大厦,邮编 100084)

[http:// www.tup.tsinghua.edu.cn](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: 31

版 次: 2000 年 1 月第 1 版 2000 年 11 月第 2 次印刷

书 号: ISBN 7-302-01093-5/TP·2210

印 数: 5001~8000

定 价: 44.00 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

**CIDR Conversion Table**

CIDR prefix- length	Dotted-Decimal	Binary	# Classfull Networks
/1	128.0.0.0	1000 0000 0000 0000 0000 0000 0000 0000	128 As
/2	192.0.0.0	1100 0000 0000 0000 0000 0000 0000 0000	64 As
/3	224.0.0.0	1110 0000 0000 0000 0000 0000 0000 0000	32 As
/4	240.0.0.0	1111 0000 0000 0000 0000 0000 0000 0000	16 As
/5	248.0.0.0	1111 1000 0000 0000 0000 0000 0000 0000	8 As
/6	252.0.0.0	1111 1100 0000 0000 0000 0000 0000 0000	4 As
/7	254.0.0.0	1111 1110 0000 0000 0000 0000 0000 0000	2 As
/8	255.0.0.0	1111 1111 0000 0000 0000 0000 0000 0000	1 A or 256 Bs
/9	255.128.0.0	1111 1111 1000 0000 0000 0000 0000 0000	128 Bs
/10	255.192.0.0	1111 1111 1100 0000 0000 0000 0000 0000	64 Bs
/11	255.224.0.0	1111 1111 1110 0000 0000 0000 0000 0000	32 Bs
/12	255.240.0.0	1111 1111 1111 0000 0000 0000 0000 0000	16 Bs
/13	255.248.0.0	1111 1111 1111 1000 0000 0000 0000 0000	8 Bs
/14	255.252.0.0	1111 1111 1111 1100 0000 0000 0000 0000	4 Bs
/15	255.254.0.0	1111 1111 1111 1110 0000 0000 0000 0000	2 Bs
/16	255.255.0.0	1111 1111 1111 1111 0000 0000 0000 0000	1 B or 256 Cs
/17	255.255.128.0	1111 1111 1111 1111 1000 0000 0000 0000	128 Cs
/18	255.255.192.0	1111 1111 1111 1111 1100 0000 0000 0000	64 Cs
/19	255.255.224.0	1111 1111 1111 1111 1110 0000 0000 0000	32 Cs
/20	255.255.240.0	1111 1111 1111 1111 1111 0000 0000 0000	16 Cs
/21	255.255.248.0	1111 1111 1111 1111 1111 1000 0000 0000	8 Cs
/22	255.255.252.0	1111 1111 1111 1111 1111 1100 0000 0000	4 Cs
/23	255.255.254.0	1111 1111 1111 1111 1111 1110 0000 0000	2 Cs
/24	255.255.255.0	1111 1111 1111 1111 1111 1111 0000 0000	1 C
/25	255.255.255.128	1111 1111 1111 1111 1111 1111 1000 0000	1/2 C
/26	255.255.255.192	1111 1111 1111 1111 1111 1111 1100 0000	1/4 C
/27	255.255.255.224	1111 1111 1111 1111 1111 1111 1110 0000	1/8 C
/28	255.255.255.240	1111 1111 1111 1111 1111 1111 1111 0000	1/16 C
/29	255.255.255.248	1111 1111 1111 1111 1111 1111 1111 1000	1/32 C
/30	255.255.255.252	1111 1111 1111 1111 1111 1111 1111 1100	1/64 C
/31	255.255.255.254	1111 1111 1111 1111 1111 1111 1111 1110	1/128 C
/32	255.255.255.255	1111 1111 1111 1111 1111 1111 1111 1111	1/256 C



# Foreword

To say that the Internet has changed in the past 10 years is a vast understatement. The Internet today is the single most influential phenomenon in computing. The way in which it continues to evolve is likely to affect every form of data communication on a global scale. Internet information-handling strategies, such as Web-based customer support and secure networked commerce, will permanently alter business and consumer behavior.

As the Internet and networking have evolved, so too has Cisco Systems' role. Cisco has become the premier worldwide provider of networking for the Internet. In this role, Cisco helps network managers build reliable, scalable, and cost-effective solutions for networks that span the requirements of small offices to globally-dispersed enterprises. Over the years, Cisco has assembled the largest and most experienced staff of networking experts who serve customers in planning, implementing, and troubleshooting large networks. Bassam Halabi is one of those experts.

Although Cisco's expertise is primarily delivered directly to customers, it also can be shared through other means. A new and important medium for sharing Cisco's networking expertise is the Cisco Press. In creating this new forum, New Riders Publishing and Cisco Systems jointly plan to build a

library of publications that provide practical information about the design and implementation of routers, switches, servers, and connectivity software.

Bassam Halabi's *Internet Routing Architectures* is the first of these books. Bassam provides his keen insight into interdomain routing, Internet service provider requirements, and the changing Internet infrastructure. He touches on issues that affect everyone connecting networks to or through the Internet. Bassam's many years of consulting experience with Cisco's largest IP-based networks contribute to this book's useful design recommendations. He brings a pragmatic perspective to the process of building links between domains and selecting network designs for specific situations.

*Internet Routing Architectures* is a readable distillation of essential design and implementation guidelines. We trust you will find it a valuable addition to your networking library.

Edward R. Kozel

*Chief Technology Officer  
Cisco Systems, Inc.*



# **Trademark Acknowledgments**

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

# Acknowledgments

---

I wish to acknowledge Ravi Chandra from the BGP4 development team for being a key reference for the protocol. I would also like to thank Robert Craig from the consulting team for his help in reviewing parts of the material; Kim Lew for motivating and leading this book project and for his guidance in organizing the material; and Laurie McGuire for her excellent input in the book development editing. Finally, this book couldn't have been written without the support and patience of my family who put up with the extra work hours during the last year.

A portion of the proceeds from this book will be donated to Cisco Networking Academies, a non-profit educational fund.



# Introduction

The Internet, an upstart academic experiment in the late 1960s, struggles with identity and success in the late 1990s. From the ARPANET to the NSF-NET to ANYBODYSNET, the Internet is no longer owned by a single entity; it is owned by anybody who can afford to buy space on it. Millions of users are seeking connectivity and thousands of companies are feeling left out if they do not tap into the Internet. This has put network designers and administrators under a lot of pressure to keep up with networking and connectivity needs. Learning about networking, and especially routing, has become a necessity.

People get surprised when networks fail and melt down; I get surprised when they don't. I say that because there is so little useful information out there. Much of the information on routing that has been available to designers and administrators up until now is doubly frustrating: The information makes you *think* that you know how to build your network until you try, and find out that you don't. I wrote *Internet Routing Architectures* to be the first book that addresses real routing issues, using real scenarios, in a comprehensive and accessible treatment.

## OBJECTIVES

The purpose of this book is to make you an expert on integrating your network into the global Internet. By presenting practical addressing, routing, and connectivity issues both conceptually and in the context of practical scenarios, the book aims to foster your understanding of routing so that you can plan and implement major network designs in an objective and informed way. Whether you are a customer or provider (or both) of Internet connectivity, this book anticipates and addresses the routing challenges facing your network.

## AUDIENCE

This book is intended for any organization that might have the need to tap into the Internet. Whether you are becoming a service provider or you are connecting to one you will find all you need to integrate your network. The perspectives of network administrators, integrators, and architects are considered throughout this book. Even though this book addresses different levels of expertise, it progresses logically from simplest to most challenging concepts and problems, and its common denominator is straightforward, practical scenarios to which anyone can relate. No major background in routing or TCP/IP is required. Any basic or background knowledge needed to understand routing is developed as needed in text discussions, rather than assumed as part of the reader's repertoire.

## ORGANIZATION

The book is organized into four parts:

- **Part 1—The Contemporary Internet.** Chapters 1–3 cover essential introductory aspects of the contemporary Internet with respect to its structure, service providers, and addressing. Even if you are already familiar with the general structure of the Internet, you are encouraged to read the portions of Chapter 1 concerning Network Access Points, the Route

Arbiter Project, and Network Information Services. The pressures that precipitated these components of the Internet have continuing practical implications for routing design problems faced by administrators. Chapter 2 provides valuable criteria by which to evaluate Internet service providers. If you represent such a provider, or are already a customer of one, some of the information may be familiar to you already.

- **Part 2—Routing Protocol Basics.** Chapters 4 and 5 cover the basics: Why interdomain routing protocols are needed and how they work. These topics are covered both generally, and in the specific context of BGP—Border Gateway Protocol—the de facto standard interdomain routing protocol of today. BGP’s particular capabilities and attributes are thoroughly introduced.
- **Part 3—Effective Internet Routing Designs.** Chapters 6–9 delve into the practical, design-oriented applications of BGP. The attributes introduced in Part 2 are shown in action, in a variety of representative network scenarios. BGP’s attributes are put to work in implementing design goals such as redundancy, symmetry, and load balancing. The challenges of making intradomain and interdomain routing work in harmony, managing growing or already large systems, and maintaining stability are addressed.
- **Part 4—Internet Routing Device Configuration.** Chapters 10 and 11 contain numerous code examples of BGP’s attributes and of various routing policies. The code examples will make the most sense to you after you have read the earlier chapters, because many of them address multiple concepts and design goals. However, so that you can juxtapose textual discussions from earlier chapters with code examples in Chapters 10 and 11, pointers entitled “Configuration Example,” have been placed in the earlier chapters. When you see one, you may wish to fast forward to the referenced page to see a configuration example of the attribute or policy being discussed.

## **APPROACH**

It is very hard to write about technical information in an accessible manner. Information that is stripped of too much technical detail loses its meaning, while complete and precise technical detail can overwhelm readers and obscure concepts. This book introduces technical detail gradually and in the context of practical scenarios whenever possible. The most heavily technical information—configuration examples in the Cisco IOS language—is withheld until the final two chapters of the book, so that it is thoroughly grounded in the concepts and sample topologies that precede it.

Although your ultimate goal is to design and implement routing strategies, it is critical to grasp concepts and principles before applying them to your particular network. This book balances conceptual and practical perspectives by following a logical, gradual progression from general to specific, and from concepts to implementation. Even in chapters and sections that necessarily take a largely descriptive approach, hands-on interests are addressed through pointers to configuration examples, frequently asked questions, and scenario-based explanation.

The scenario-based approach is an especially important component of this book: It utilizes representative network topologies as a basis for illustrating almost every protocol attribute and routing policy discussed. Even though you may not see your exact network situation illustrated, the scenario is specific enough to facilitate learning-by-example, and general enough that you can extrapolate how the concepts illustrated apply to your situation.

## **FEATURES AND TEXT CONVENTIONS**

This book works hard not to withhold protocol details and design-oriented information, while at the same time recognizing that building general and conceptual

understanding necessarily comes first. Two features are included to help emphasize what is practical and design-oriented as underlying concepts are developed:

- **Pointers to configuration examples**—located in the margins next to pertinent text discussions, these references point forward to places in Chapters 10 and 11 where related configuration examples can be found.
- **Frequently Asked Questions**—located at the end of every chapter, these questions anticipate practical and design-oriented questions you may have, for your particular network, after having read the chapter.

In addition, this book utilizes several other conventions:

- **Troubleshooting margin notes**—these brief references are placed adjacent to text discussions that are especially relevant to preventing or correcting common routing problems and programming mistakes.
- **Notes**—these set-off passages include elaborations that will further illuminate text passages, but which can be skipped without loss of understanding of core topics.

# Contents at a Glance

---

## **PART 1      THE CONTEMPORARY INTERNET 1**

### **Chapter 1      Evolution of the Internet 3**

Origins of the Internet 3

The Internet Today 6

NSFNET Solicitations 7

Network Access Points 8

Route Arbiter Project 13

The very high-speed Backbone Network Service (vBNS) 17

Moving the Regional Providers 19

NSF Solicits NIS Managers 20

Other Internet Registries 29

Internetworking Routing Registries (IRR) 30

Looking Ahead 31

Frequently Asked Questions 33

References 34

### **Chapter 2      ISP Services and Characteristics 37**

ISP Services 38

ISP Pricing and Technical Characteristics 39

Looking Ahead 48  
Frequently Asked Questions 49

**Chapter 3    Handling IP Address Depletion 53**

Overview of IPv4 Addressing 54  
Basic Addressing 54  
IP Address Space Depletion 60  
Looking Ahead 82  
Frequently Asked Questions 83  
References 85

**PART 2       ROUTING PROTOCOL BASICS 87**

**Chapter 4    Interdomain Routing Basics 89**

Overview of Routers and Routing 89  
Segregating the World into Administrations 95  
Border Gateway Protocol Version 4 101  
UPDATE Message and Routing Information 113  
Looking Ahead 119  
Frequently Asked Questions 120  
References 122

**Chapter 5    Tuning BGP Capabilities 125**

Building Peer Sessions 125  
Sources of Routing Updates 132  
Overlapping Protocols: Backdoors 140  
The Routing Process Simplified 141  
Controlling BGP Routes 147  
Route Filtering and Attribute Manipulation 169  
BGP4 Aggregation 180  
Looking Ahead 186  
Frequently Asked Questions 187  
References 190

<b>PART 3</b>	<b>EFFECTIVE INTERNET ROUTING DESIGNS 191</b>
<b>Chapter 6</b>	<b>Redundancy, Symmetry, and Load Balancing 193</b>
	Redundancy 194
	Symmetry 201
	Load Balancing 201
	Specific Scenarios: Designing Redundancy, Symmetry, and Load Balancing 203
	Looking Ahead 228
	Frequently Asked Questions 229
<b>Chapter 7</b>	<b>Controlling Routing Inside the Autonomous System 233</b>
	Interaction of Non-BGP Routers with BGP Routers 233
	BGP Policies Conflicting with Internal Defaults 236
	Policy Routing 244
	Looking Ahead 248
	Frequently Asked Questions 249
<b>Chapter 8</b>	<b>Controlling Large-Scale Autonomous Systems 253</b>
	Route Reflectors 253
	Confederations 263
	Controlling IGP Expansion 267
	Virtual Private Networks with Route Reflectors 275
	Looking Ahead 277
	Frequently Asked Questions 279
	References 280
<b>Chapter 9</b>	<b>Designing Stable Internets 283</b>
	Route Instabilities on the Internet 283
	BGP Stability Features 288
	Looking Ahead 294
	Frequently Asked Questions 295