


TN915 04/1



CCIE Professional Development.

Routing TCP/IP, Volume I

TCP/IP 路由技术, 卷 I

Jeff Doyle



清华大学出版社

(京)新登字 158 号

CCIE Professional Development: Routing TCP/IP, Volume I

Jeff Doyle

“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: 01-99-0591

图书在版编目(CIP)数据

TCP/IP 路由技术 第 1 卷:英文/多依尔(Doyle,J.)著. - 影印版. - 北京:清华大学出版社, 1999.4

(CISCO 系列丛书)

ISBN 7-302-02949-0

I . T… II . 多… III . 因特网-路由选择-英文 IV . TP393.4

中国版本图书馆 CIP 数据核字(1998)第 09340 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: 66

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

书 号: ISBN 7-302-02949-0/TP·1560

印 数: 3001~7000

定 价: 89.00 元



Foreword

In today's world of networking, mission-critical networks are being designed for data, voice, and video. Due to different traffic patterns and the quality of service required by each type of information, solid hands-on experience is imperative for managing, designing, and troubleshooting these networks.

Attaining a strong degree of hands-on experience translates into in-depth understanding of the concepts, scalability, and deployment issues of today's networks. Such experience also builds the expertise to analyze traffic patterns and the knowledge of when, where, and how to apply protocol and bandwidth features to enhance performance.

To help further your hands-on experience, Cisco Press is publishing the CCIE Professional Development series of books. Books in this series will significantly help your understanding of protocol concepts, and they provide real-world examples and case studies to strengthen the theoretical concepts examined. I highly recommend that you use these books as a hands-on learning tool by duplicating the examples and case studies using Cisco products. You can even take this further by tweaking the configuration parameters to see which changes each network goes through by using the extensive debugging features provided in each Cisco product.

In the first book of the CCIE Professional Development series, *CCIE Professional Development: Routing TCP/IP, Volume I*, Jeff Doyle does a fantastic job of building

the TCP/IP concepts, from IP address classes to analyzing protocol metrics. Each chapter contains examples, network topologies with IP addresses, packet analysis, and Cisco debugging outputs. In my opinion, the best parts are the case studies, in which Jeff compares different features of the protocol by using more or less the similar topology. This generates a strong understanding of the protocol concepts and features.

I recommend *CCIE Professional Development: Routing TCP/IP, Volume I* for any networking certification, and I believe that it also makes an excellent university networking course book.

Imran Qureshi
CCIE Program Manager

PART I


Routing Basics

Chapter 1—Basic Concepts: Internetworks, Routers, and Addresses

Chapter 2—TCP/IP Review

Chapter 3—Static Routing

Chapter 4—Dynamic Routing Protocols

- 
- **Bicycles with Motors**
 - **Data Link Addresses**
 - **Repeaters and Bridges**
 - **Routers**
 - **Network Addresses**

Basic Concepts: Internetworks, Routers, and Addresses

Once upon a time, computing power and data storage were centralized. Mainframes were locked away in climate-controlled, highly secure rooms, watched over by a priesthood of IS administrators. Contact with a computer was typically accomplished by bringing a stack of Hollerith cards to the priests, who interceded on our behalf with the Big Kahuna.

The advent of the minicomputer took the computers out of the IS temple of corporations and universities and brought them to the departmental level. For a mere \$100K or two, engineering and accounting and any other department with a need for data processing could have their own machines.

Following on the heels of the minicomputers were microcomputers, bringing data processing right to the desktop. Affordability and accessibility dropped from the departmental level to the individual level, making the phrase *personal computer* part of everyone's vocabulary.

Desktop computing has evolved at a mind-boggling pace, but it was certainly not an immediate alternative to centralized, mainframe-based computing. There was a ramping-up period in which both software and hardware had to be developed to a level where personal computers could be taken seriously.

BICYCLES WITH MOTORS

One of the difficulties of decentralized computing is that it isolates users from one another and from the data and applications they may need to use in common. When a file is created, how is it shared with Tom, Dick, and Harriet down the hall? The early solution to this was the storied SneakerNet: Put the file on floppy disks and hand carry them to the necessary destinations. But what happens when Tom, Dick, and Harriet modify their copies of the file? How does one ensure that all information in all versions are synchronized? What if those three coworkers are on different floors or in different buildings or cities? What if the file needs to be updated several times a day? What if there are not three coworkers, but 300 people? What if all 300 people occasionally need to print a hard copy of some modification they have made to the file?

The *local-area network*, or LAN, is a small step back to centralization. LANs are a means of pooling and sharing resources. Servers enable everyone to access a common copy of a file or a common database; no more “walkabouts” with floppies, no more worries about inconsistent information. E-mail furnishes a compromise between phone calls, which require the presence of the recipient, and physical mail service, which is called snail mail for a good reason. The sharing of printers and modem pools eliminates the need for expensive, periodically used services on every desk.

Of course, in their infancy, LANs met with more than a little derision from the mainframe manufacturers. A commonly heard jibe during the early years was, “A LAN is like a bike with a motor, and we don’t make Mopeds!” What a difference a few years and a few billion dollars would make.

Physically, a LAN accomplishes resource pooling among a group of devices by connecting them to a common, shared medium, or

data link. This medium may be twisted-pair wires (shielded or unshielded), coaxial cable, optical fiber, infrared light, or whatever. What matters is that all devices attach commonly to the data link through some sort of network interface.

Data link

A shared physical medium is not enough. Rules must govern how the data link is shared. As in any community, a set of rules is necessary to keep life orderly, to ensure that all parties behave themselves, and to guarantee that everyone gets a fair share of the available resources. For a local-area network, this set of rules, or *protocol*, is generally called a *Media Access Control* (MAC). The MAC, as the name implies, dictates how each machine will access and share a given medium.

So far, a LAN has been defined as being a community of devices such as PCs, printers, and servers coexisting on a common communications medium and following a common protocol that regulates how they access the medium. But there is one last requirement: As in any community, each individual must be uniquely identifiable.

DATA LINK ADDRESSES

In a certain community in Colorado, two individuals are named Jeff Doyle. One Jeff Doyle frequently receives telephone calls for the person with whom he shares a name—so much so that his clever wife has posted the correct number next to the phone to redirect errant callers to their desired destination. In other words, because two individuals cannot be uniquely identified, data is occasionally delivered incorrectly and a process must be implemented to correct the error.

Among family, friends, and associates, a given name is usually sufficient for accurately distinguishing individuals. However, as this example shows, most names become inaccurate over a larger

population. A more unique identifier, such as a United States Social Security number, is needed to distinguish one person from every other.

Frame

Devices on a LAN must also be uniquely and individually identified or they, like humans sharing the same name, will receive data not intended for them. When data is to be delivered on a LAN, it is encapsulated within an entity called a *frame*, a kind of binary envelope. Think of data encapsulation as being the digital equivalent of placing a letter inside an envelope, as in Figure 1.1¹. A destination address and a return (source) address are written on the outside of the envelope. Without a destination address, the postal service would have no idea where to deliver the letter. Likewise, when a frame is placed on a data link, all devices attached to the link “see” the frame; therefore, some mechanism must indicate which device should pick up the frame and read the enclosed data.

Figure 1.2 shows the format of most common LAN frames. Notice that every case includes a destination address and a source address. The format of the address depends on the particular MAC protocol, but all the addresses serve the same purpose: to uniquely identify the machine for which the frame is destined and the device from which it was sent.

The three most common data links currently used in LANs are Ethernet, Token Ring, and FDDI. Although each link is drastically different from the others, they share a common format for addressing devices on the network. This format, originally standardized by Xerox’s Palo Alto Research Center (PARC)² and now administered by the Institute of Electrical and Electronics Engineers

As will be seen later, creating a data link layer frame is really more like putting an envelope inside a larger envelope.

The full name, as reading any modern text on networking will tell you, is The Now Famous Xerox PARC.




Figure 1.1

Encapsulation means putting data into a frame—a kind of digital “envelope” for delivery.

(IEEE), is variously called the burned-in address,³ the physical address, the machine address, or most commonly, the MAC address.

The MAC address is a 48-bit number, which, as Figure 1.3 shows, is designed so that every device anywhere on the planet should be uniquely identifiable. Most everyone has heard the legends of large batches of network interface cards being turned out with identical burned-in addresses by unscrupulous “cloning” companies or as the result of “stuck” programming code. Although most of those stories are nothing more than legends, one can imagine what would happen if all devices on a LAN had the same MAC address: Imagine a town in which every resident is named Wess-
vick Smackley. Men, women, children, dogs, and cats all named

³ The address is usually permanently programmed, or burned in, to a ROM on the network interface.

Ethernet

PREAMBLE	DESTINATION ADDRESS	SOURCE ADDRESS	TYPE	DATA	FRAME CHECK SEQUENCE
----------	---------------------	----------------	------	------	----------------------

IEEE 802.3

PREAMBLE	DESTINATION ADDRESS	SOURCE ADDRESS	LENGTH	DATA	FRAME CHECK SEQUENCE
----------	---------------------	----------------	--------	------	----------------------

IEEE 802.5/TOKEN RING

S D	A C	F C	DESTINATION ADDRESS	SOURCE ADDRESS	DATA	FRAME CHECK SEQUENCE	E D
--------	--------	--------	---------------------	----------------	------	----------------------	--------

FDDI

PREAMBLE	S D	F C	DESTINATION ADDRESS	SOURCE ADDRESS	DATA	FRAME CHECK SEQUENCE	E D	F S
----------	--------	--------	---------------------	----------------	------	----------------------	--------	--------

SD = Start Delimiter

AC = Access Control

FC = Frame Control

ED = End Delimiter

FS = Frame Status

Figure 1.2

The frame format of a few common LAN data link frames.

Wessvick Smackley. Everyday communication, not to mention the career of the town gossip, would be unimaginably difficult.⁴

Although the MAC addresses are by convention referred to as “addresses,” they are really names. Think about it: Because the

⁴ In real life, duplicate MAC addresses on a network are most likely to occur as the result of network administrators using locally administered addresses. This occurrence is common enough on Token Ring networks that one step of the Token Ring insertion process is a duplicate address check.




Figure 1.3

A MAC address.

identifier is burned in, or permanently assigned, to a device, it is a part of that device and goes wherever the device goes.⁵

Most adults have several street addresses through their lives, but few have more than one given name. A name identifies an entity—whether a person or a PC. An address describes where that person or PC is located.

In the interest of clarity, this book uses the term *data link identifier* or *MAC identifier* instead of MAC address. The reason for making such a distinction will soon be clear.

REPEATERS AND BRIDGES

The information presented so far may be distilled into a few brief statements:

- A data communication network is a group of two or more devices connected by a common, shared medium.

⁵ Although some data link addresses may be or must be administratively configured, the point is that they are identifiers, unique within a network.

- These devices have an agreed-upon set of rules, usually called the Media Access Control, or MAC, that govern how the media is shared.
- Each and every device has an identifier, and each identifier is unique to only one device.
- Using these identifiers, the devices communicate by encapsulating the data they need to send within a virtual envelope called a *frame*.

So here's this wonderful resource-sharing tool called a LAN. It's so wonderful, in fact, that everyone wants to be connected to it. And herein is the rub. As a LAN grows, new problems present themselves.

The first problem is one of physical distance. Figure 1.4 shows that three factors can influence an electrical signal. These factors may decrease or eliminate any intelligence the signal represents:

- Attenuation
- Interference
- Distortion

As the distance the signal must travel down the wire increases, so do the degrading effects of these three factors. Photonic pulses traveling along an optical fiber are much less susceptible to interference but will still succumb to attenuation and distortion.

Repeaters are added to the wire at certain intervals to alleviate the difficulties associated with excessive distance. A repeater is placed on the media some distance from the signal source but still near enough to be able to correctly interpret the signal (see Figure 1.5). It then repeats the signal by producing a new, clean copy of the old degraded signal. Hence, the name *repeater*.




Figure 1.4

Attenuation, interference, and distortion prevent a signal from arriving in the same shape it was in when it left. Attenuation (a) is a function of the resistance of the wire. A certain amount of signal energy must be spent “pushing past” the resistance. Interference (b) is a function of outside influences—noise—which adds characteristics to the signal that should not be there. Distortion (c) is a function of the wire impeding different frequency components of the signal in different ways.




Figure 1.5

By placing a repeater in the link at a distance where the original signal can still be recognized, despite the effects of attenuation, interference, and distortion, a fresh signal can be generated and the length of the wire extended.

A repeater may be thought of as part of the physical medium. It has no real intelligence, but merely regenerates a signal; a digital

repeater is sometimes facetiously called a “bit spitter” for this reason.

The second problem associated with growing LANs is congestion. Repeaters are added to extend the distance of the wire and to add devices; however, the fundamental reason for having a LAN is to share resources. When a too-large population tries to share limited resources, the rules of polite behavior begin to be violated and conflicts erupt. Among humans, poverty, crime, and warfare may result. On Ethernet networks, collisions deplete the available bandwidth. On Token Ring and FDDI networks, the token rotation time and timing jitter may become prohibitively high.

Drawing boundaries between populations of LAN devices is a solution to overcrowding. This task is accomplished by the use of *bridges*.⁶

Figure 1.6 shows the most common type of bridge: a *transparent bridge*. It performs three simple functions: learning, forwarding, and filtering. It is transparent in that end stations have no knowledge of the bridge.

The bridge learns by listening *promiscuously* on all its ports. That is, every time a station transmits a frame, the bridge examines the source identifier of the frame. It then records the identifier in a *bridging table*, along with the port on which it was heard. The bridge therefore learns which stations are out port 1, which are out port 2, and so on.

In Figure 1.6, the bridge uses the information in its bridging table to forward frames when a member of one population—say, a sta-

⁶ If you cut through the marketing hype surrounding modern Ethernet and Token Ring switches, you'll find that these very useful tools are merely high-performance bridges.

tion out port 1—wants to send a frame to a member of another population: a station out port 2.




Figure 1.6

The transparent bridge segments network devices into manageable populations. A bridging table tracks the members of each population and manages communication between the populations.

A bridge that only learns and forwards would have no use. The real utility of a bridge is in the third function, filtering. Figure 1.6 shows that if a station out port 2 sends a frame to another station out port 2, the bridge will examine the frame. The bridge consults its bridging table and sees that the destination device is out the same port on which the frame was received and will not forward the frame. The frame is filtered.

Bridges enable the addition of far more devices to a network than would be possible if all the devices were in a single population,