



PEARSON NEW INTERNATIONAL EDITION

Internetworking with TCP/IP
Volume One
Douglas E. Comer
Sixth Edition

Pearson New International Edition

Internetworking with TCP/IP
Volume One
Douglas E. Comer
Sixth Edition

PEARSON®

Pearson Education Limited

Edinburgh Gate
Harlow
Essex CM20 2JE
England and Associated Companies throughout the world

Visit us on the World Wide Web at: www.pearsoned.co.uk

© Pearson Education Limited 2014

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a licence permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

All trademarks used herein are the property of their respective owners. The use of any trademark in this text does not vest in the author or publisher any trademark ownership rights in such trademarks, nor does the use of such trademarks imply any affiliation with or endorsement of this book by such owners.

PEARSON®

ISBN 10: 1-292-04081-5
ISBN 13: 978-1-292-04081-3

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

Printed in the United States of America

Table of Contents

Chapter 1 Introduction And Overview Douglas E. Comer	1
Chapter 2 Overview Of Underlying Network Technologies Douglas E. Comer	19
Chapter 3 Internetworking Concept And Architectural Model Douglas E. Comer	35
Chapter 4 Protocol Layering Douglas E. Comer	47
Chapter 5 Internet Addressing Douglas E. Comer	69
Chapter 6 Mapping Internet Addresses To Physical Addresses (ARP) Douglas E. Comer	101
Chapter 7 Internet Protocol: Connectionless Datagram Delivery (IPv4, IPv6) Douglas E. Comer	119
Chapter 8 Internet Protocol: Forwarding IP Datagrams Douglas E. Comer	147
Chapter 9 Internet Protocol: Error And Control Messages (ICMP) Douglas E. Comer	167
Chapter 10 User Datagram Protocol (UDP) Douglas E. Comer	185
Chapter 11 Reliable Stream Transport Service (TCP) Douglas E. Comer	199
Chapter 12 Routing Architecture: Cores, Peers, And Algorithms Douglas E. Comer	247
Chapter 13 Routing Among Autonomous Systems (BGP) Douglas E. Comer	263
Chapter 14 Routing Within An Autonomous System (RIP, RIPng, OSPF, IS-IS) Douglas E. Comer	289

Chapter 15 Internet Multicasting	319
Douglas E. Comer	
Chapter 16 Label Switching, Flows, And MPLS	353
Douglas E. Comer	
Chapter 17 Packet Classification	369
Douglas E. Comer	
Chapter 18 Mobility And Mobile IP	381
Douglas E. Comer	
Chapter 19 Network Virtualization: VPNs, NATs, And Overlays	399
Douglas E. Comer	
Chapter 20 Client-Server Model Of Interaction	419
Douglas E. Comer	
Chapter 21 The Socket API	431
Douglas E. Comer	
Chapter 22 Bootstrap And Autoconfiguration (DHCP, NDP, IPv6-ND)	463
Douglas E. Comer	
Chapter 23 The Domain Name System (DNS)	485
Douglas E. Comer	
Chapter 24 Electronic Mail (SMTP, POP, IMAP, MIME)	511
Douglas E. Comer	
Chapter 25 World Wide Web (HTTP)	525
Douglas E. Comer	
Chapter 26 Voice And Video Over IP (RTP, RSVP, QoS)	539
Douglas E. Comer	
Chapter 27 Network Management (SNMP)	559
Douglas E. Comer	
Chapter 28 Software Defined Networking (SDN, OpenFlow)	583
Douglas E. Comer	
Chapter 29 Internet Security And Firewall Design (IPsec, SSL)	605
Douglas E. Comer	
Index	623

Chapter Contents

- 1.1 The Motivation For Internetworking, 1
- 1.2 The TCP/IP Internet, 2
- 1.3 Internet Services, 2
- 1.4 History And Scope Of The Internet, 6
- 1.5 The Internet Architecture Board, 7
- 1.6 The IAB Reorganization, 8
- 1.7 Internet Request For Comments (RFCs), 8
- 1.8 Internet Growth, 9
- 1.9 Transition To IPv6, 12
- 1.10 Committee Design And The New Version of IP, 12
- 1.11 Relationship Between IPv4 And IPv6, 13
- 1.12 IPv6 Migration, 14
- 1.13 Dual Stack Systems, 15
- 1.14 Organization Of The Text, 15
- 1.15 Summary, 16

1

Introduction And Overview

1.1 The Motivation For Internetworking

Internet communication has become a fundamental part of life. Social networks, such as Facebook, provide connections among a group of friends and allow them to share interests. The World Wide Web contains information about such diverse subjects as politics, atmospheric conditions, stock prices, crop production, and airline fares. Family and friends use the Internet to share photos and keep in touch with VoIP telephone calls and live video chats. Consumers use the Internet to purchase goods and services and for personal banking. Companies take orders and make payments electronically. The move to cloud computing will put more information and services online.

Although it appears to operate as a unified network, the Internet is not engineered from a single networking technology because no technology suffices for all uses. Instead, networking hardware is designed for specific situations and budgets. Some groups need high-speed wired networks to connect computers in a single building. Others need a low-cost wireless network for a private home. Because low-cost hardware that works well inside a building cannot span large geographic distances, an alternative must be used to connect sites that are thousands of miles apart.

In the 1970s, a technology was created that makes it possible to interconnect many disparate individual networks and operate them as a coordinated unit. Known as *internetworking*, the technology forms the basis for the Internet by accommodating multiple, diverse underlying hardware technologies, providing a way to interconnect the networks, and defining a set of communication conventions that the networks use to interoperate. The internet technology hides the details of network hardware, and permits computers to communicate independent of their physical network connections.

Internet technology is an example of *open system interconnection*. It is called *open* because, unlike proprietary communication systems available from one specific vendor, the specifications are publicly available. Thus, any individual or company can build the hardware and software needed to communicate across the Internet. More important, the entire technology has been designed to foster communication among machines with diverse hardware architectures, to use almost any packet switched network hardware, to accommodate a wide variety of applications, and to accommodate arbitrary computer operating systems.

1.2 The TCP/IP Internet

In the 1970s and 1980s, U.S. government agencies realized the importance and potential of internet technology, and funded research that made possible a global Internet[†]. This book discusses principles and ideas that resulted from research funded by the *Defense Advanced Research Projects Agency (DARPA)*[‡]. The DARPA technology includes a set of network standards that specify the details of how computers communicate, as well as a set of conventions for interconnecting networks and forwarding traffic. Officially named the *TCP/IP Internet Protocol Suite* and commonly referred to as *TCP/IP* (after the names of its two main standards), it can be used to communicate across any set of interconnected networks. For example, TCP/IP can be used to interconnect a set of networks within a single building, within a physical campus, or among a set of campuses.

Although the TCP/IP technology is noteworthy by itself, it is especially interesting because its viability has been demonstrated on a large scale. It forms the base technology for the global Internet that connects approximately two billion individuals in homes, schools, corporations, and governments in virtually all populated areas of the planet. An outstanding success, the Internet demonstrates the viability of the TCP/IP technology and shows how it can accommodate a wide variety of underlying hardware technologies.

1.3 Internet Services

One cannot appreciate the technical details underlying TCP/IP without understanding the services it provides. This section reviews internet services briefly, highlighting the services most users access, and leaves to later chapters the discussion of how computers connect to a TCP/IP internet and how the functionality is implemented.

Much of our discussion of services will focus on standards called *protocols*. Protocol specifications, such as those for TCP and IP, define the syntactic and semantic rules for communication. They give the details of message formats, describe how a computer responds when a message arrives, and specify how a computer handles errors or other abnormal conditions. Most important, protocols allow us to discuss computer communication independent of any particular vendor's network hardware. In a sense, protocols

[†]We will follow the usual convention of capitalizing *Internet* when referring specifically to the global Internet, and use lower case to refer to private internets that use TCP/IP technology.

[‡]At various times, *DARPA* has been called the *Advanced Research Projects Agency (ARPA)*.

are to communication what algorithms are to computation. An algorithm allows one to specify or understand a computation without knowing the details of a particular programming language or CPU instruction set. Similarly, a communication protocol allows one to specify or understand data communication without depending on detailed knowledge of a particular vendor's network hardware.

Hiding the low-level details of communication helps improve productivity in several ways. First, because they can use higher-level protocol abstractions, programmers do not need to learn or remember as many details about a given hardware configuration. Thus, they can create new network applications quickly. Second, because software built using higher-level abstractions are not restricted to a particular computer architecture or a particular network hardware, the applications do not need to be changed when computers or networks are replaced or reconfigured. Third, because applications built using higher-level protocols are independent of the underlying hardware, they can be ported to arbitrary computers. That is, a programmer does not need to build a special version of an application for each type of computer or each type of network. Instead, applications that use high-level abstractions are more general-purpose — the same code can be compiled and run on an arbitrary computer.

We will see that the details of each service available on the Internet are given by a separate protocol. The next sections refer to protocols that specify some of the application-level services as well as those used to define network-level services. Later chapters explain each of the protocols in detail.

1.3.1 Application Level Internet Services

From a user's point of view, the Internet appears to consist of a set of application programs that use the underlying network to carry out useful tasks. We use the term *interoperability* to refer to the ability of diverse computing systems to cooperate in solving computational problems. Because the Internet was designed to accommodate heterogeneous networks and computers, interoperability was a key requirement. Consequently, Internet application programs usually exhibit a high degree of interoperability. In fact, most users access applications without understanding the types of computers or networks being used, the communication protocols, or even the path data travels from its source to its destination. Thus, a user might access a web page from a desktop system connected to a cable modem or from an iPad connected to a 4G wireless network.

The most popular and widespread Internet application services include:

- *World Wide Web*. The Web became the largest source of traffic on the global Internet between 1994 and 1995, and remains so. Many popular services, including Internet search (e.g., Google) and social networking (e.g., Facebook), use web technology. One estimate attributes approximately one quarter of all Internet traffic to Facebook. Although users distinguish among various web-based services, we will see that they all use the same application-level protocol.

- *Cloud Access And Remote Desktop.* Cloud computing places computation and storage facilities in *cloud data centers*, and arranges for users to access the services over the Internet. One access technology, known as a *remote desktop service*, allows a user to access a computer in a remote data center as if the computer is local. The user only needs an interface device with a screen, keyboard, mouse or touchpad, and a network connection. When the data center computer updates the video display, the remote desktop service captures the information, sends it across the Internet, and displays it on the user's screen. When the user moves the mouse or presses a key, the remote desktop service sends the information to the data center. Thus, the user has full access to a powerful PC, but only needs to carry a basic interface device such as a tablet.
- *File Transfer.* The file transfer protocol allows users to send or receive a copy of a data file. Many file downloads, including movie downloads, invoke a file transfer mechanism. Because they often invoke file transfer from a web page, users may not be aware that a file transfer application has run.
- *Electronic Mail (email).* Electronic mail, which once accounted for large amounts of Internet traffic, has largely been replaced by web applications. Many users now access email through a web application that allows a user to read messages in their mailbox, select a message for processing, and forward the message or send a reply. Once a user specifies sending a message, the underlying system uses an email transfer protocol to send the message to the recipient's mailbox.
- *Voice And Video Services.* Both streaming video and audio already account for a nontrivial fraction of bits transported across the global Internet, and the trend will continue. More important, a significant change is occurring; video upload is increasing, especially because users are using mobile devices to send video of live events.

We will return to a discussion of applications in later chapters and examine them in more detail. We will see exactly how applications use the underlying TCP/IP protocols, and why having standards for application protocols has helped ensure that they are widespread.

1.3.2 Network-Level Internet Services

A programmer who creates network applications has an entirely different view of the Internet than a user who merely runs applications such as web browsers. At the network level, the Internet provides two broad services that all application programs use. While it is unimportant at this time to understand the details of the services, they are fundamental to an overview of TCP/IP:

- *Connectionless Packet Delivery Service.* Packet delivery, explained in detail throughout the text, forms the basis for all internet services. Connectionless delivery is an abstraction of the service that most packet-switching networks offer. It means simply that a TCP/IP internet forwards small messages from one computer to another based on address information carried in the message. Because it