

ROUTING IN TODAY'S **Internetworks**

The Routing Protocols of
IP, DECnet, NetWare, and AppleTalk



Mark Dickie

VNR Communications Library




30802923

802923

Routing in Today's Internetworks

The Routing Protocols of IP, DECnet, NetWare, and AppleTalk



Mark Dickie



VAN NOSTRAND REINHOLD
New York

Copyright © 1994 by Van Nostrand Reinhold

Library of Congress Catalog Card Number

ISBN 0-442-01811-8

All Rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without the written permission of the publisher.

I T P Van Nostrand Reinhold is an International Thomson Publishing company.
ITP logo is a trademark under license.

Printed in the United States of America

Van Nostrand Reinhold
115 Fifth Avenue
New York, NY 10003

International Thomson Publishing Germany
Königswinterer Str. 418
53227 Bonn
Germany

International Thomson Publishing
Berkshire House, 168-173
High Holborn, London WC1V 7AA
England

International Thomson Publishing Asia
221 Henderson Building #05-01
Singapore 0315

Thomas Nelson Australia
102 Dodds Street
South Melbourne 3205
Victoria, Australia

International Thomson Publishing Japan
Kyowa Building, 3F
2-2-1 Hirakawacho
Chiyoda-Ku, Tokyo 102
Japan

Nelson Canada
1120 Birchmount Road
Scarborough, Ontario
M1K 5G4, Canada

ARCFF 16 15 14 13 12 11 10 9 8 7 6 5 4 3

Library of Congress Cataloging in Publication Data

Dickie, Mark.

Routing in today's internetworks : the routing protocols of IP, DECnet,
NetWare, and AppleTalk / Mark Dickie.

p. cm.

Includes index.

ISBN 0-442-01811-8

1. Computer network protocols. 2. Computer network architectures.

I. Title.

TK5105.5.D52 1993

004.6'7--dc20

93-33790

CIP

Preface

In today's work environment the network is becoming more and more of a standard work tool than ever before, and as this "tool" continues to evolve it continues to become more complex. Companies continue to add more computing resources to their environment, and as this trend continues, the need to tie together computer networks becomes of greater and greater importance. Today's networking environment has become more complex, and this complexity is growing every day. This makes the job of the network designer, the network manager, and the network administrator more difficult. The need to link various distributed networks within a company or between companies into a single cohesive user network has driven the creation of the internetworking marketplace, and the need for devices such as bridges, routers, and gateways.

Many books are written about the data communications world but few if any have had a focus on the world of routing in today's internetworks. The goal of this book is to provide the network manager, network administrator, and network designer with information on the latest movement in the routing world. Routing protocols have been around for some time, and they are responsible for allowing computer users to get from one network to another. The routing world is undergoing a change, moving away from the older routing protocols such as Routing Information Protocol (RIP), DECnet Routing Protocol (DRP), Internet Packet Exchange Routing Information Protocol (IPX RIP), and Routing Table Maintenance Protocol (RTMP). Newer routing protocols such as Open Shortest Path First (OSPF), Intermediate System–Intermediate System (IS–IS), Netware Link Services Protocol (NLSP), and Apple Update-based Routing Protocol (AURP) are the latest efforts by the network community to increase the efficiency of how routers communicate with each other.

This book discusses four networking environments, Internet Protocol (IP), DECnet, Novell NetWare, and AppleTalk. As we explore each protocol group we will first review some basics of the particular group, describe the current routing methodology in use today, and then explore the new routing protocol that is being introduced into this particular routing

environment. Where possible we will discuss any relevant implementation issues to consider when creating or converting to one of these new protocols.

Before we dig into the various protocol suites we will start with a general discussion of internetwork technology. A review of the OSI Model will lead off to ensure everyone is on the right footing. Following this discussion will be sections on bridging (including Spanning Tree Bridge and Source Route Bridging), generic routing technology, and gateways. This will give all readers an understanding of the basics that form the underpinning for the protocol discussions that follow later in the book.

In the section on the Internet Protocol we will cover a variety of topics. We will begin with a discussion of the generic IP world. This will include discussions of the addressing structure, subnetwork addressing, address resolution, and the datagram format. We will then discuss the current IP routing environment, which will cover the Routing Information Protocol (RIP), and how RIP is implemented in most of the internet community to allow data flow across the Internet. Routing table maintenance will be discussed, and this will include a discussion of the various problems found in today's RIP-based networks.

Once this review of IP is concluded, we will focus on the Open Shortest Path First (OSPF) protocol, which was created to allow more flexibility and functionality in the implementation of IP networks. This discussion includes how OSPF is used to allow the creation and maintenance of the IP routing environment, the advantages that OSPF offers over RIP, and any relevant implementation issues. Examples will be used throughout the discussion to help clarify technical points. The IP section of the book will contain many examples of various basic routing principles. These principles will hold true for many of the other protocols, and the examples will not be repeated in following sections. You should read the chapters in the IP section first to make later discussions about other protocols more meaningful.

In the DECnet section we will discuss the DECnet Phase IV environment. This will include, as in the IP discussion, the address structure, address resolution, and various packet formats. This discussion will also cover the concept of areas and node types. We will then cover the current DECnet Phase IV routing implementation using the DECnet Routing Protocol (DRP). A look at the DECnet Phase V implementation, which is based on the OSI Intermediate System-Intermediate System (IS-IS), will follow. Digital is the first major vendor to implement the OSI routing protocols, so it is appropriate to use this environment for our discussion.

The Novell section begins with a discussion of the generic Novell network architecture. We will discuss the Internetwork Packet Exchange (IPX), its addressing scheme, packet format, and address resolution. Next we will discuss Novell's implementation of RIP and its use in maintaining routing information for Novell networks. Once today's Novell environment is described we can discuss the next iteration of the Novell network, in which Novell Link Services Protocol (NLSP) is used to provide routing across Novell internetworks. Examples will be used as needed to help clarify technical points.

In the AppleTalk section a discussion of AppleTalk Phase 2 will set the stage for the following discussion. A review of the Datagram Delivery Protocol (DDP) will give us a base from which to work. Apple's addressing scheme, address resolution, zones, packet formats, and name binding will be discussed. Following this will be a discussion of the Routing Table Maintenance Protocol (RTMP), and how this Apple implementation of RIP is used in today's AppleTalk networks to maintain the routing of data through an AppleTalk internetwork. The Apple Update-based Routing Protocol (AURP) will be the next topic discussed to show how Apple is moving forward to give its users a more robust routing environment for the creation of complex AppleTalk networks.

The only major convention used in this book that the reader must be aware of is how field lengths are stated in packet format diagrams. All field lengths are stated in number of bits, not bytes. In reality, many fields are less than one byte in length so I decided to keep everything in one unit of measure instead of confusing the reader with bits in one place and bytes in another place.

The other major goal of this book is to discuss this somewhat complicated subject matter in an easy-to-understand manner. My hope is to take the mystery out of a technology that is not understood by most of the networking world, yet is a technology that is becoming increasingly more important in the design of today's internetworks. I will try not to waste words, or to simply retype the various protocol specifications to you. All the information discussed in this book is available in the public domain through various organizations. Unfortunately, most of these specifications were not written for the average person who is involved in the networking environment. Over the last 2.5 years as a technical instructor at Wellfleet Communications I have been teaching these same routing concepts to individuals implementing the technology in their corporate networks. Today, in my job as a Principal Instructor I attempt to relate these subjects in a way that is easy to understand, and I will try to bring this style into the writing of this book.

Finally, I would like to thank all those people that have helped me in the writing of this book. I have to thank my wife Chris and my daughter Jessica for putting up with me working on this book at all hours of the day, including our family vacation. Several of my peers, especially Eural Authemont and Dennis Baker, have helped in the technical editing of the book to ensure that I have not said anything blatantly stupid, and some individuals have made very pertinent suggestions on how to make this book more enjoyable. For those of you interested in discussing this book with me, or for those of you who would like to make suggestions for future revisions of the book, I can be reached through the Internet at mail address MDICKIE@WELLFLEET.COM.

Contents

Preface	xi
Chapter 1. So What Is Routing	1
The OSI Reference Model	2
Layer 1—The Physical Layer	3
Layer 2—The Data Link Layer	3
Layer 3—The Network Layer	4
Layer 4—The Transport Layer	5
Layer 5—The Session Layer	5
Layer 6—The Presentation Layer	6
Layer 7—The Application Layer	6
Bridge Concepts	7
Spanning Tree Bridge	10
Source Route Bridge	12
Routing Concepts	12
Gateway Concepts	19
Summary	19
Chapter 2. Internet Protocol (IP)	21
IP Addressing	21
Subnet Addressing	23
Address Resolution	28
IP Packet	30
The Future of IP	34
Summary	35
Chapter 3. Routing Information Protocol (RIP)	37
Distance Vector Algorithm	37
Building a Routing Table	39
Routing Updates	42
Convergence of the Internetwork	44
RIP Limitations	49
Summary	50
Chapter 4. Open Shortest Path First (OSPF)	51
Shortest Path First Protocol	52
Topological Database	54
Shortest Path Tree	54
Dijkstra's Algorithm	56
Areas	61

The Backbone	64
Routing Between Areas	65
Router Classifications	65
Network Types	66
Designated Routers	67
External Routes	68
Stub Areas	68
Type of Service Routing	69
IP Variable-Length Subnetting	70
Equal Cost Multipath	72
Authentication	73
OSPF Paths and Costs	73
Summary	74
Chapter 5. OSPF Protocol	77
OSPF Message Handling	77
OSPF Data Structures	80
Forming Adjacencies	81
Link State Advertisements	90
Router Links Advertisement	91
Network Links Advertisement	92
Summary Links Advertisement	92
Autonomous System External Links Advertisement	93
Flooding Procedure	96
Routing Table Maintenance	100
OSPF Implementation Issues	103
Summary	106
Chapter 6. DECnet Phase IV	107
Digital Network Architecture Overview	107
DNA Physical and Data-Link Layers	110
DECnet Routing Protocol	111
Addressing	112
Areas	113
Costs and Paths	114
Designated Router	116
Decision Process	120
Update Process	121
Forwarding Process	124
Receive Process	125
Routing Table	125
Summary	127
Chapter 7. DECnet Phase V	129
DECnet Phase V Architecture	129
Connectionless-mode Network Services (CLNS)	131
Addressing	131
CLNS Protocol Functions	132
Address Resolution	137

End System Hello Message Fields.....	138
Intermediate System Hello Message Fields.....	140
Redirect Message Fields	141
CLNS Packet Format.....	143
Summary	148
Chapter 8. Intermediate System-Intermediate System (IS-IS)	149
System Types.....	150
Subnetwork Types.....	150
Functions/Processes	151
Addresses.....	152
Designated Routing Nodes	153
Decision Process	154
Update Process.....	157
Forwarding Process.....	165
Receive Process.....	166
Summary	167
Chapter 9. Novell Netware	167
NetWare Protocol Stack.....	167
Client-Server Interaction.....	168
Data-Link Layer Protocols.....	170
Internetwork Packet Exchange (IPX).....	171
Service Advertising Protocol (SAP).....	175
NetWare Core Protocol (NCP).....	179
NetBIOS.....	180
Summary	181
Chapter 10. IPX Routing Information Protocol (RIP)	183
RIP Message Format.....	183
RIP Operation	184
Split Horizon Algorithm	186
Router Initialization	187
Router Shutdown	188
Router Information Maintenance	188
Summary	189
Chapter 11. NetWare Link Services Protocol (NLSP)	191
NLSP General Characteristics	191
NLSP WAN Support	193
NLSP Adjacencies	200
NLSP and RIP/SAP.....	205
Differences Between IS-IS and NLSP.....	208
Summary	210
Chapter 12. AppleTalk	211
AppleTalk Protocol Suite	211
Datagram Delivery Protocol (DDP).....	213
Sockets	213
AppleTalk Address Structure.....	215
Address Acquisition	215

Address Resolution.....	217
DDP Message Format	217
Name Binding Protocol	219
Zone Information Protocol (ZIP).....	223
Summary	228
Chapter 13. Routing Table Maintenance Protocol (RTMP)	231
Types of AppleTalk Routers	231
AppleTalk Routing Model	233
Routing Tables and Maintenance	234
RTMP Message Formats.....	236
Summary	239
Chapter 14. AppleTalk Update-based Routing Protocol	241
AppleTalk Tunneling	242
AURP Information Exchange.....	247
AURP Model	247
Initial Information Exchange	249
Routing Updates.....	250
Network Hiding.....	252
Resolving Network Number Conflicts	252
Hop Count Reduction.....	254
Routing Loops.....	255
Alternate Paths	255
Summary	256
Glossary	257
Bibliography	265
Index	267

So What Is Routing?

Today's network environment is a complex one. Users at all levels of the organization are involved in the use and administration of the network, various departments are in charge of purchasing equipment that is going onto networks, and everyone is trying to figure out how to allow one user to talk to another user. Everyone is looking for a way to consolidate hardware and software resources to make more efficient use of the working environment, while still providing to the individual user the tools necessary to produce at the highest level.

This need to allow the sharing of resources, both hardware and software, drove the growth of the networking market in the 1980s. As the use of Local Area Networks (LAN) and Wide Area Networks (WAN) grew so did the need to connect users in different departments, different offices, and different companies. This demand for distributed sharing of resources was the impetus for the creation of the internetworking market. This internetworking market spawned many companies producing a variety of hardware and software products, most notably bridges, routers, and gateways.

This chapter covers some networking basics to give everyone a sound base of understanding. It is intended for the network novice who is just getting into the routing world, but the more experienced network user may also find this a good refresher. The topics to be discussed in this chapter will include:

- The Open Systems Interconnection (OSI) Reference Model
- An overview of Bridge Concepts
- An overview of Routing Concepts
- An overview of Gateway Concepts

This discussion will set the stage for the more detailed discussion of routing protocols. It is important to understand where routing fits in the network scheme, and that is the goal of this chapter.

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

Figure 1-1 OSI Reference Model

THE OSI REFERENCE MODEL

During the 1980s network usage began to grow at a phenomenal rate. Most vendors were creating proprietary network architectures that forced customers to buy everything from a single vendor. As network usage grew, so did the need to integrate network devices from various vendors. Due to the proprietary nature of most vendors' products, this integration was a very difficult task. As time went on, pressure from the user community grew for the creation of a single, universally adopted network architecture that would allow true multivendor interoperability.

The International Organization for Standardization (ISO) took on the task of creating such a network architecture. Through the efforts of the ISO came something known as the Open Systems Interconnection (OSI) Reference Model. This OSI Model was an attempt to define the various services that a network was supposed to provide to its users. If you sit and think about all the tasks that a network is supposed to provide you, the user, the list will become very long, very quickly. However, the OSI working group made an attempt to do just that, define, in a general sense, all the services that a network should provide to its users. This was no simple task as you can imagine, but they did it, and once all the services were identified they were grouped together into logical groupings known as layers. A *layer* is a group of related services that a network is to provide for the users. The OSI Model has seven layers defined, these layers are shown in Figure 1-1.

Let's take a look at the basic functions of each of the seven layers in the model, and then how the model would be used to allow communication between two stations on a network.

Layer 1—The Physical Layer

Starting at the bottom of the model, we will work our way up through each of the layers. The Physical Layer deals with the physical connection to the network. At this layer we describe the various mechanical, electrical, and functional specifications for connecting a device to a particular type of network. The type of connector, the size of the connector, the number of pins in the connector, the function of each of the pins, the acceptable electrical range for the signals on the network, and the type of cable used are some of the major items defined in a particular specification for a network. Other items that are defined at the Physical Layer are network speed and encoding method (how to recognize a binary 1 or 0). It is these definitions that allow a network station to physically attach to a specific network.

Examples of Physical Layer specifications are 10Base5, 10Base2, and 10Base-T. These are all specifications for various types of LAN networks. Examples of WAN specifications would be EIA 232-D, 449, or CCITT V.35.

Layer 2—The Data Link Layer

The Data-Link Layer provides a variety of very important functions, and is actually looked at as two different sublayers by most organizations. The two sublayers are the Media Access Control (MAC), and the Logical Link Control (LLC).

The functions of the MAC are to define how we get around on a particular network. The MAC defines items such as frame format (the standard unit of measure at the Data-Link Layer), addressing, flow control, synchronization, and media access. The Data-Link Layer uses a message known as a *frame* in order to transport data from one point to another. Each type of network has a specific frame format associated with it. This defines the standard "envelope" for carrying data from one point to another on that network. The address defined at this layer is known by any one of the following names; hardware address, MAC address, station address, physical address, or Ethernet address. No matter what name you decide to use, each typically refers to a 48-bit address that defines a station's location on a particular network. Because of my background I have come to use the term MAC address to describe this address, and will use this term throughout the remainder of the book. This MAC address is similar to a street address—your house is identified on your street by a number so that the mailman can find it, and a network station is given a MAC address so

that it can be located on that particular network. Every station on the network must have a unique MAC address so that it can be identified without confusion. Every network also defines its own media access rules. These can be thought of as the rules of the road for a network. When you are in your car driving and you encounter a red light, you stop the car because you know that rules of the road say that it is no longer your turn to use the roads; it is someone else's turn. In networking the same concept applies. Each type of network defines its own rules of the road which dictate how a particular station may access the network. Examples of these access methods are Carrier Sense/Multiple Access with Collision Detection (CSMA/CD) and token passing.

The Logical Link Control (LLC) provides a more abstract service. The best way I know to describe this service is to think about a mailman carrying a letter. The letter is in an envelope with an address on it; the address identifies that the letter is going to 159 Main Street. The mailman finds 159 Main Street, but when he sees the building he notices that it is a fifteen-story apartment tower. Upon entering the lobby of the building the mailman is faced with three hundred mailboxes. The mailman looks back down at the envelope to discover that the author of the letter did add the appropriate box number to the address, and the mailman now makes final delivery of the letter by placing it in the appropriate mailbox. The LLC provides something of the same service in which the network is the street, the MAC address is the street address of the station, and the mailbox number is something called a *socket*, or a Service Access Point (SAP). A computer on a network could be providing a variety of functions. Various users may wish to access the network through this one particular device, so that they may send and receive data on the network. Data arrives at a station based on its MAC address (street address), but final delivery is made to the appropriate user based on the socket number (mailbox number). This allows two processes on the same network to maintain a "logical connection" with each other for the purpose of data exchange.

The Data-Link Layer provides the functions that allow users to bridge networks together. We will discuss the concepts of transparent bridging, spanning tree bridging, and source route bridging later in this chapter.

Layer 3— The Network Layer

The Network Layer is where the network provides its routing functionality. It is important to understand the differences between Layers 2 and 3. At Layer 2 we deal with MAC addresses, which we equated to street addresses. By looking at the street address alone, it is not possible to

determine with any kind of certainty where that street is in relation to you, or how to get to that street in the most efficient manner? Layer 2 MAC addresses allow you to get from one point on a particular network to another point on the same network. At Layer 3 however, we form an address that allows us to go from one point on one network to another point on a different network. It is the job of the router to help you get your data from one network to another in the most efficient way available.

At Layer 3 we define a variety of items such as packet format (Layer 3 uses packets to transfer data as opposed to Layer 2's frames), addressing, address resolution, and routing table maintenance. Most Layer 3 schemes define an addressing scheme that allows a station and its particular network to be identified. Address resolution deals with mapping these Layer 3 addresses to the appropriate Layer 2 addresses. We are going to spend a lot of time talking about Layer 3 and the various addressing schemes used, but when data needs to get from one point to another on the same network it is a MAC address that must be used to accomplish this task. Layer 3 also defines the method for creating and maintaining the routing environment. This is the main purpose of this book, to discuss the new methods being employed in various well-known Layer 3 architectures for maintaining this routing environment.

Layer 4—The Transport Layer

The Transport Layer provides the means for maintaining reliable data transmission from end to end across a network. All the error checking for data, end-to-end acknowledgments, retransmissions, and data sequencing are handled by the Transport Layer. There is the ability to do some error checking down at Layer 2, however this error checking is referred to as bit-level error checking and does not constitute data-level error checking. What that means is that the frame check sequence that can be employed at Layer 2 simply tells one station that it received the same set of bits that was sent to it by the transmitting station. This error check does not guarantee that the bits received are of any particular value to the receiver. It is the duty of Layer 4 to provide for the reliable transfer of data from one point to another, end to end.

Layer 5—The Session Layer

Layer 5 provides the ability for two processes on two different stations to form a connection for the purpose of communication with one another. This is often referred to as *dialogue management* and is employed by protocols

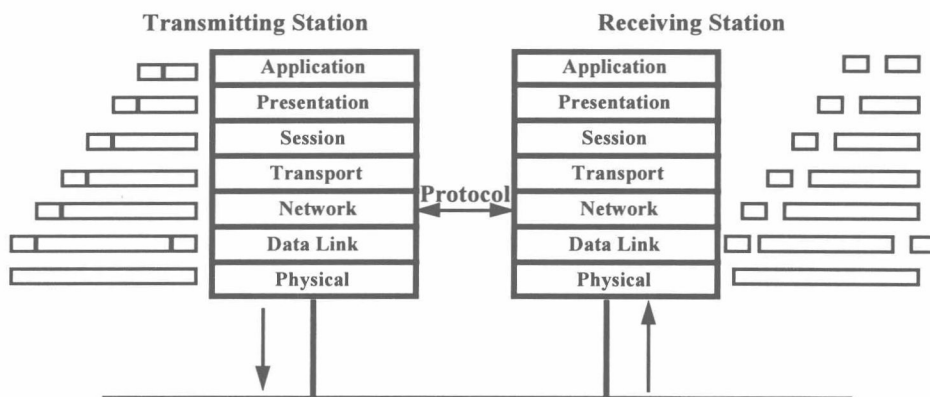


Figure 1-2 OSI Communications

such as NetBIOS. Layer 5 provides the means for the session between two processes to be established and removed, and is also responsible for the synchronization between the two processes in the session.

Layer 6—The Presentation Layer

The Presentation Layer is responsible for ensuring that the data received is in the proper format. Data conversion, syntax conversion, or protocol conversion are within the realm of Layer 6. It is at Layer 6 that gateways are provided for the purpose of protocol conversion. More on this in a little while.

Layer 7—The Application Layer

The Application Layer provides the user with access to the network environment. This is where a variety of network applications are written to grant users the ability to perform a variety of functions on the network.

When this model is put together it allows stations to communicate with each other. If you look closely, you will find that each layer provides a completely separate set of functions to the network, but all the layers work together as part of a team to make the network environment function. Let's take a quick look at how the model would be used to allow communication between two stations.

Figure 1-2 shows how data is passed through the OSI model as it is transmitted and received. As a piece of data is transferred from one station to another it begins its journey at Layer 7 of the transmitting station. The data passes down from layer to layer through the model, each layer adding