




INTERNETWORKING CONCEPTS GUIDE

<http://www.3com.com/>

Part No. 980-000076/001
Published: September 1997



3Com Corporation ■ 5400 Bayfront Plaza ■ Santa Clara, California ■ 95052-8145

© **3Com Europe Limited, 1997.** All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without permission from 3Com Europe Limited.

3Com Europe Limited reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Europe Limited to provide notification of such revision or change.

3Com Europe Limited provides this documentation without warranty of any kind, either implied or expressed, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

UNITED STATES GOVERNMENT LEGENDS:

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following restricted rights:

For units of the Department of Defense:

Restricted Rights Legend: Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) for Restricted Rights in Technical Data and Computer Software Clause at 48 C.F.R. 52.227-7013. 3Com Europe Limited, c/o Merchants' House, Wilkinson Road, Cirencester, Gloucestershire, GL7 1YT United Kingdom.

For civilian agencies:

Restricted Rights Legend: Use, reproduction or disclosure is subject to restrictions set forth in subparagraph (a) through (d) of the Commercial Computer Software – Restricted Rights Clause at 48 C.F.R. 52.227-19 and the limitations set forth in 3Com Corporation's standard commercial agreement for the software. Unpublished rights reserved under the copyright laws of the United States.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hard-copy documentation, or on the removable media in a directory file named LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, AccessBuilder, Boundary Routing, EtherLink, NETBuilder, OfficeConnect and SuperStack are registered trademarks of 3Com Corporation. ServiceConnect is a trademark of 3Com Corporation.

AppleTalk, AppleShare, EtherTalk, LaserWriter, LocalTalk, Macintosh, and TokenTalk are registered trademarks of Apple Computer, Inc. Novell, NetWare and Yes NetWare are registered trademarks of Novell Inc. Windows, WIndows 95 and the Windows logo are registered trademarks of Microsoft Corporation. VT100 is a registered trademark of Digital Equipment Corporation. UNIX is a registered trademark, licensed exclusively through X/Open Company Ltd.

Other brand and product names may be registered trademarks or trademarks of their respective holders.

Environmental Statement:

It is 3Com's policy to be environmentally friendly in all its operations. This manual is printed on paper that comes from sustainable, managed European forests. The production process for making the pulp has a reduced AOX level (adsorbable organic halogen) resulting in elemental chlorine-free paper.

This paper is fully biodegradable and recyclable.

CONTENTS

ABOUT THIS GUIDE

Finding Specific Information in This Guide	1
Conventions	2
Related Documentation	2

1 INTERNETWORKING OVERVIEW

Introduction	1-1
The OSI Reference Model	1-1
Application Layer	1-2
Presentation Layer	1-2
Session Layer	1-3
Transport Layer	1-3
Network Layer	1-3
Data Link Layer	1-3
Physical Layer	1-3
Protocols	1-4
Addressing	1-6
Link Layer (MAC) Addresses	1-6
Network Layer Addresses	1-6
Internetworking Devices	1-6
Gateways and Hosts	1-7
Routers	1-8
Bridges	1-8
Repeaters	1-8

2 THE INTERNET PROTOCOLS

Introduction	2-1
The Internet Protocol Suite	2-1
Internet Protocol (IP)	2-1
Transmission Control Protocol (TCP)	2-2

User Datagram Protocol (UDP)	2-2
Other Protocols	2-2
Routing Information Protocol (RIP)	2-2
Serial Line Internet Protocol (SLIP)	2-3
Point-to-Point Protocol (PPP)	2-3
Simple Mail Transport Protocol (SMTP)	2-3
Simple Network Management Protocol (SNMP)	2-4
File Transfer Protocol (FTP)	2-4
Telnet	2-4
IP Addressing	2-4
IP Address Classes	2-5
Class A	2-5
Class B	2-5
Class C	2-5
Class D	2-6
Class E	2-6
IP Address Notation	2-6
Subnetting	2-8
Subnet Masks	2-9
Worked Example	2-10
Define the Subnet Mask	2-10
Assign the Host Address	2-11

3 NETWARE PROTOCOLS

NetWare Internetworking Protocols	3-1
The IPX Protocol	3-1
IPX Addressing	3-1
Novell Routing Information Protocol (NRIP)	3-2
Service Advertisement Protocol (SAP)	3-3

4 APPLE TALK

Introduction	4-1
About AppleTalk	4-1
AppleTalk Addressing	4-2
Network Nodes	4-2
Network Numbers	4-2
Zone Names	4-3

AppleTalk Named Objects	4-3
AppleTalk Routers	4-4
Seed and Non-Seed Mode	4-4

5 BRIDGING INTERNETWORKS

Introduction	5-1
Bridging Versus Routing	5-1
Bridging Concepts	5-2
Bridging and the OSI Reference Model	5-2
Transparent Bridging	5-3
Learning	5-3
Filtering	5-4
Forwarding	5-5
Active Loops	5-5
Broadcast Storms	5-5
Incorrect Learning of MAC Addresses	5-6
Spanning Tree	5-7
Spanning Tree Problems	5-7
Local and Remote Bridging	5-8
Advantages and Disadvantages of Bridging	5-8
Advantages	5-8
Disadvantages	5-9

6 ROUTING INTERNETWORKS

Introduction	6-1
Routing Concepts	6-1
Routing Tables	6-2
Static Routes	6-2
Switching	6-3
Routing and the OSI Reference Model	6-3
Bridge/Router	6-4
Routing Protocols	6-4
IPX Routing	6-5
IP Routing	6-5
Advantages	6-5
Disadvantages	6-6

7 REMOTE ACCESS USING ISDN

- Introduction 7-1
- How ISDN Works 7-3
 - ISDN Logical Channels 7-3
 - ISDN User Interface Standards 7-4
- ISDN for Remote Access 7-5
 - Primary Link Backup 7-6
 - Dial on Congestion 7-6
 - Dial-on-Demand 7-7
 - Telecommuting 7-8
- Security Management 7-9
- Minimizing Costs 7-9

GLOSSARY

INDEX

ABOUT THIS GUIDE

About This Guide provides an overview of this guide, describes guide conventions, tells you where to look for specific information, and lists other publications that may be useful.

This guide describes some of the basic principles behind common internetworking technologies.

This guide is intended for internetworking novices and those who wish to improve their internetworking knowledge.

Finding Specific Information in This Guide

This table shows the location of specific information in this guide.

If you are looking for:	Turn to:
An introduction to the principles of internetworking.	Chapter 1
An introduction to Internet Protocols.	Chapter 2
An introduction to Novell NetWare®.	Chapter 3
An introduction to AppleTalk®.	Chapter 4
An introduction to bridges and bridging.	Chapter 5
An introduction to routers and routing.	Chapter 6
An introduction to ISDN.	Chapter 7
A glossary of internetworking terms.	Glossary

Conventions

Table 1 and Table 2 list conventions that are used throughout this guide.

Table 1 Notice Icons




Icon	Notice Type	Description
	Information note	Important features or instructions
	Caution	Information to alert you to potential damage to a program, system, or device
	Warning	Information to alert you to potential personal injury


Table 2 Text Conventions

Convention	Description
Words in <i>italicized</i> type	Italics emphasize a point or denote new terms at the place where they are defined in the text.
Words in bold-face type	Bold text denotes key features.

Related Documentation

This document is intended as additional background and preparatory information for the following documents from each document set:

- *Getting Connected Guide.*
- *Software Reference.*



INTERNETWORKING OVERVIEW

This chapter gives a basic overview of the internetworking environment. The information provided here is intended as a foundation to the remainder of the chapters in this guide.

Introduction

One of the most challenging tasks in the computer industry is moving information between computers of diverse design. This linking of computer systems, software, and communications devices into strategic infrastructures is called internetworking.

In an effort to standardize the various protocols and make the networking implementations of different vendors interoperable, the International Standards Organization (ISO) developed the Open Systems Interconnection (OSI) reference model. Although other models have been proposed, it is the OSI reference model that has become the industry standard to describe how data communications take place.

The OSI Reference Model

The OSI Reference Model provides a common basis for the coordination of standards development for systems interconnection, whilst allowing existing standards to be placed into perspective within its structure.

The OSI model separates the functions required for effective computer communications (such as error-checking and addressing) into seven 'layers'. This was agreed by the organizations involved as an appropriate number for achieving a manageable analysis of the functions involved in data communication.

Each layer sends packets of information to the layers above and below it, but each layer only *understands* information that comes from the same layer on another stack.

The layers are numbered from one to seven. Layer seven is the layer closest to the user and layer one can be considered the layer closest to the computer hardware. This layer structure is illustrated in the table below and then described in greater detail.

Table 1-1 The Seven Layer of the OSI Reference Model

Layer	Function	Description
7	Application	Selects appropriate service for applications (user interface).
6	Presentation	Provides code conversion and data reformatting.
5	Session	Co-ordinates interaction between end-to-end application processes.
4	Transport	Provides end-to-end data integrity and quality of service.
3	Network	Switches and routes information to the appropriate network device.
2	Data Link	Transfers units of information to other end of the physical link.
1	Physical	Performs transmission/reception on the network medium.

Application Layer

The application layer is the layer closest to the user. It provides information services to support the application processes which reside outside of the OSI model.

Presentation Layer

The presentation layer formats the data which is presented to the application layer. It ensures that data from the application layer of one system is readable by the application layer of another system. In simple terms, it can be viewed as a translator of information.

Session Layer

The session layer allows two applications to synchronize and manage their data exchange. It sets up a communication channel between two application or presentation layers for the duration of the network transaction, manages the communication, and terminates the connection. This is known as a session.

Transport Layer

The transport layer is the interface between the layers concerned with application issues, and those concerned with data transport issues. It provides the session layer with reliable message transfer facilities. It also offers transparent transfer of data between end stations, error recovery, and flow control. You could say that it provides a transparent pipe for the interchange of information, supporting whatever level of reliability is appropriate for the application.

Network Layer

The network layer controls the operation of the network or subnetwork. It decides which physical path the data should take based on such factors as network conditions and priorities of service. It establishes, maintains and terminates connections between end-systems, taking care of all addressing, routing, and facility selection.

Data Link Layer

The data link layer provides reliable transmission of data across a physical link. By providing error control, it allows the network layer to assume error free data transmission.

Physical Layer

The physical layer handles the electrical and mechanical interface to the communications media. This includes procedures for activating, maintaining, and de-activating the physical connection. It is responsible for converting data from the layers above it into electrical signals compatible with the communications media.

Protocols

Protocols are sets of rules that define how different parts of a network interact to allow devices to communicate with one another. They provide a ‘common language’ to allow different vendors’ computer equipment to communicate with each other. The different devices can use completely different software, provided that each device’s software can agree on the meaning of the data.

Protocols can describe low-level details of machine-to-machine interfaces (like the order in which bits and bytes are sent across the wire) or high-level exchanges between allocation programs (the way in which two programs transfer a file across the Internet). Various protocols work at different layers of the seven-layer OSI reference model (Figure 1-1).

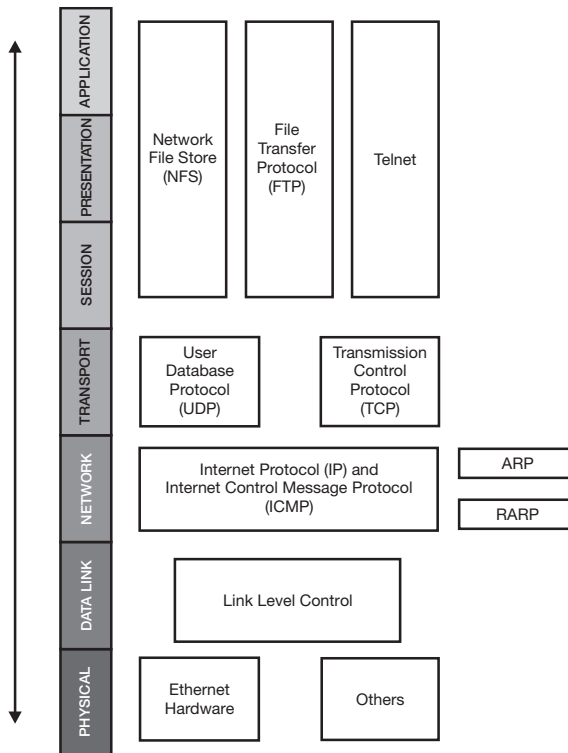


Figure 1-1 Protocols and the OSI Reference Model

Usually several protocols operate simultaneously to give full network functionality. In most cases, multiple protocols in an internetworking environment are related to one another as members of what is known as a *protocol stack* (sometimes called a protocol suite).

Data to be transported between two stations on a network is split into manageable blocks called *packets* (frames in bridging terminology). In addition to the data being transferred, each packet contains control information used for error checking, addressing, and other purposes.

The content of the control information is defined by the network protocols used. Often *multiple protocols* co-exist within a single packet with each protocol defining a different part of the packet control information.

When multiple protocols are used, the protocol control information is appended to the data in sequential order corresponding with the OSI reference model. The highest layer protocol first, then each subsequent protocol in the protocol stack. This process is called *enveloping* (Figure 1-2).

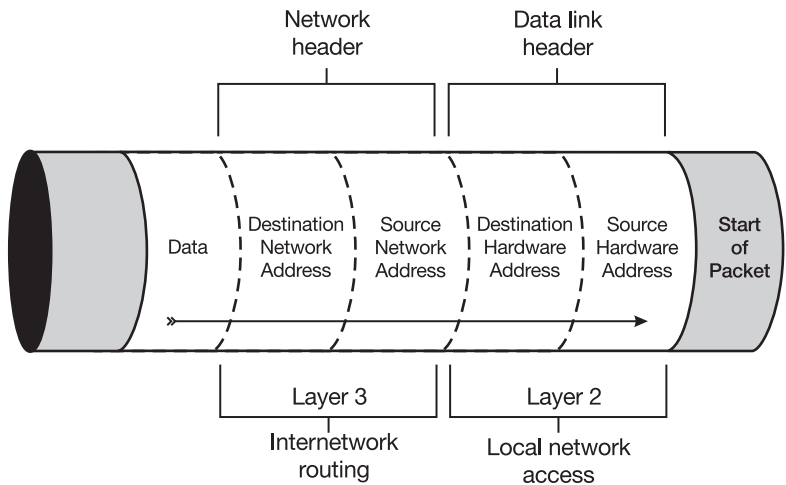


Figure 1-2 Enveloping

The enveloping pattern illustrated here is common in the communications industry. However, the tasks assigned to each protocol in the packet differ for different vendor's implementations.

Addressing

Addressing is a vital part of internetworking technology. The location of each device on a network must be uniquely identified in order for information to be directed to it. The two main types of network address are described in the following paragraphs.

Link Layer (MAC) Addresses

Link layer addresses are also known as *physical*, *hardware* or *MAC* addresses. They are usually unique for each networked device.

Link layer addresses exist at layer two of the OSI reference model. Most networked devices have only one physical network connection and thus have only one link layer address.

Network Layer Addresses

Network layer addresses are also known as *logical* or *virtual* addresses. They are usually in a hierarchical format—like a postal address. This means they can be sorted as they go along, because each line of the address narrows the search.

Network layer addresses usually exist at layer three of the OSI reference model. Their format varies according to the protocol used.

Internetworking Devices

The OSI Reference Model provides a simple representation of how information moves through a network. It can serve as a basis for understanding and characterizing an overall networking strategy. The relationship of the various internetworking devices to the OSI Reference Model is shown in [Figure 1-3](#).

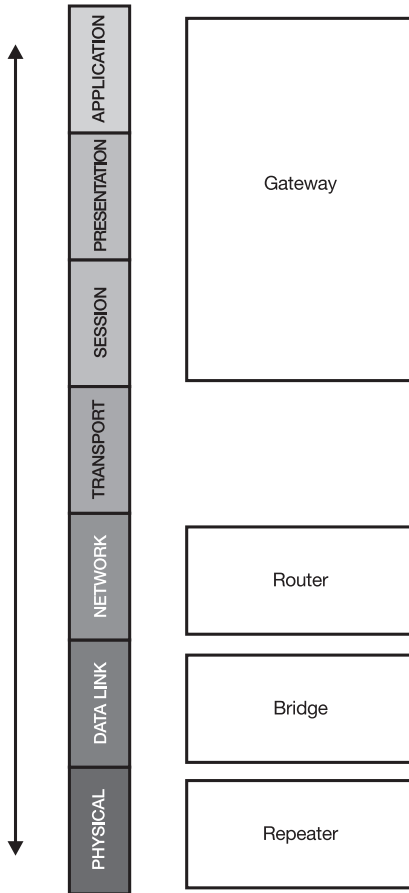


Figure 1-3 Internetworking Devices and the OSI Reference Model

Gateways and Hosts

Gateways and hosts operate at the Session, Presentation, and Application layers. They connect different networking environments such as SNA and DECnet protocols. Gateways normally relate to specific applications and network configurations, and they may use a protocol converter to translate data from one set of protocols to another. Hosts are generally less complex in functionality and may simply be a networked PC workstation.

Routers

Routers usually operate at the Network layer (they may sometimes operate as part of the transport layer too). They connect networks into internetworks that are physically unified, but in which each network retains its identity as a separate network environment. A router's primary purpose is to find the best path from one network environment to another and forward packets between them.

Bridges

Bridges operate at the Data Link layer. They usually connect similar network environments into logical and physical single internetworks. Latterly translation bridges have been developed to connect dissimilar LAN types. Bridges store and forward data in frames, and are transparent end-to-end stations.

Repeaters

Repeaters operate at the Physical layer. They receive transmissions (bits) on a Local Area Network (LAN) segment and regenerate the bits to boost a degraded signal and to extend the length of the LAN segment. They are not technically internetworking devices because they only extend to a single logical LAN segment, but they are typically spoken of as one.



THE INTERNET PROTOCOLS

This chapter provides a basic introduction to the Internet protocols.

Introduction

The Internet protocols can be used to communicate across any set of 'like-minded' interconnected networks. They are equally well suited for local area network (LAN) and wide area network (WAN) communications. They are vendor independent and can support multiple technologies. The standards documents are called RFCs and are written and maintained by the Internet Engineering Task Force (IETF). Copies of the RFCs can be found on the IETF website which is currently located at the following URL:

<http://www.ietf.cnri.reston.va.us/>

The Internet Protocol Suite

The Internet protocol suite consists of a well-defined set of communications protocols and several standard application protocols. Transmission Control Protocol/Internet Protocol (TCP/IP) is probably the most widely known and is a combination of two of the protocols (IP and TCP) working together. TCP/IP is an internationally adopted and supported networking standard which provides connectivity between equipment from many vendors over a wide variety of networking technologies.

Internet Protocol (IP)

The Internet Protocol defines a connectionless data delivery service between networked devices. Packets of data are sent as datagrams across the network. Large packets may be broken into several datagrams which are each sent individually across the network. Each datagram carries its full destination address and control information.

It is routed through the network independent of all other datagrams. No connections or logical circuits are established between the devices that are communicating.

A datagram consists of a header and a data segment. The header contains routing and processing information. The data segment contains the actual data to be transferred.

Transmission Control Protocol (TCP)

The Transmission Control Protocol works with IP to provide reliable delivery. It ensures that the various datagrams which make up a single packet of information are reassembled in the correct order at their destination address. It also ensures that missing datagrams are resent until they are received intact.

The primary purpose of TCP is to avoid the loss, damage, duplication, delay, or misordering of packets that can occur under IP. When IP forwards datagrams, there is no guarantee that the datagrams will arrive. If they do arrive, they will not necessarily be in the correct order. TCP adds reliability to IP. It also provides security mechanisms.

User Datagram Protocol (UDP)

The User Datagram Protocol is an alternative to TCP. It also provides data transfer, but without many of the reliable delivery capabilities of TCP. UDP is faster than TCP because it has fewer security features, and is useful when guaranteed data delivery is not of paramount importance.

Other Protocols

In addition to the lower layer protocols described above, the suite contains numerous other protocols that support applications such as file transfer, electronic mail, network management, and remote login. Some common IP protocols are described below.

Routing Information Protocol (RIP)

To route packets in an internetwork, IP uses a dynamic routing protocol called the Routing Information Protocol (RIP). Today RIP is the most commonly used Interior Gateway Protocol (IGP) in the

Internet community. The primary function of RIP is to inform routers about other routers on the network.

Different protocols use differing network characteristics or *metrics* when making routing decisions. The metric commonly employed by RIP is a *Hop Count*. A hop count is defined by the number of routing nodes there are between the source and destination units.

Approximately every 30 seconds, each IP router will advertise to all other routers on the internetwork how many hops it takes to reach all connected logical networks. This count is based on the router's network position and the state of its physical links. In this way each router has up-to-date information about the state of the network enabling it to make (and assist other routers to make) decisions about the best routes to use for data transmission.

Serial Line Internet Protocol (SLIP)

SLIP transmits IP packets over serial lines. If using SLIP, the network must use TCP/IP as its primary means of communication between resources. A SLIP connection only allows one communication application to be active at any one time.

Point-to-Point Protocol (PPP)

PPP also transmits packets over serial point-to-point links. It is one of the most popular methods for dial up connections to the Internet, because it allows other standard protocols to be used (such as TCP/IP and Novell IPX) over standard telephone connections. PPP can also be used for LAN connections. It supports multiple communications applications and is widely used with ISDN links.

PPP does introduce an additional connection time overhead. It is also more complex to configure in networked devices as much of the additional information it uses is unique to the connecting service.

Simple Mail Transport Protocol (SMTP)

SMTP transfers e-mail from one server to another across the network. End users must use the Post Office Protocol (POP) to transfer the messages to their machines.

Simple Network Management Protocol (SNMP)

SNMP is used to manage nodes and/or devices on an IP network. It provides a means to monitor and set network configurations and runtime parameters. It may also be used to gather statistical information about network performance.

File Transfer Protocol (FTP)

FTP provides a way to move files between computer systems. It is a widely used way of transferring files to and from the Internet and is relatively simple to operate.

Telnet

Telnet is the Internet standard terminal-emulation protocol for connecting to remote terminals. When Telnet is used to connect to a remote device, the user can use that remote machine as if it were local to them.

IP Addressing

Each device (or host) on the Internet is assigned a unique address. These devices/hosts may be personal computers, communications servers, ports on a communications server, internetwork routers, network control servers, or UNIX machines.

Some devices, such as routers, have physical connections to more than one network, and these must normally be assigned a unique internet address for *each* network connection. The internet then behaves like a virtual network, using these assigned addresses when sending or receiving packets of information.

Each internet address has a 32-bit address field. This field is split into two parts: the first part identifies the network on which the host resides, and the second part identifies the host itself. Thus hosts attached to the same network share a common prefix designating their network number.

IP Address Classes

There are five classes of IP address. Each begins with a unique bit pattern, which is used by the Internet software residing on network hosts to identify the address class. Once the internet software has identified the address class, it can determine which bits represent the network number and which bits determine the host portion of the address.

Any of the address classes can be used in a private TCP/IP network, providing that connections outside of that private network (to other TCP/IP networks) are never needed.

If a private IP addressing number scheme is established within a private corporate network, connections out of that network to external public or other private TCP/IP networks can be achieved via a computer which has software enabling it to act as an IP gateway. This will, if configured correctly, provide the IP numbering/address translation between the networks.

All registered IP addresses are assigned by InterNIC. The InterNIC website is currently located at the following URL:

<http://ds.internic.net/>

Class A

A Class A network address has the highest order bit set to zero, a seven-bit network number, and a 24-bit local host address. Class A addressing can specify up to 126 networks with up to 16,777,214 hosts per network.

Class B

A Class B network address has the two highest order bits set to 1-0, a 14-bit network number, and a 16-bit local host address. Class B addressing can specify up to 16,382 networks with up to 65,534 hosts per network.

Class C

A Class C network address has the three highest order bits set to 1-1-0, a 21-bit network number, and an 8-bit local host address.

Class C addressing can specify up to 2,097,152 networks with up to 254 hosts per network.

Class D

Class D is reserved for multicast addresses with the four highest order bits set to 1-1-1-0. They are usually used to identify a group of network devices which all run a common application program or networking software.

Class E

An IP address with the highest order bits set to 1-1-1-1-0 is reserved for future use as a Class E IP address.

IP Address Notation

For simplicity, IP addresses are specified as four decimal numbers each separated by a dot. This format is called dotted decimal notation.

The notation divides the 32-bit IP address into four 8-bit (byte) fields called octets, and specifies the value of each field independently as a decimal number with the fields separated by dots.

The addresses are given in the form <network><host> when viewed in binary notation. A mask determines how much of the address belongs to the network, and how much to the host.

For a standard Class B address, the mask is 255.255.0.0. The first two octets refer to the network and the last two are available for hosts (PCs, servers, and the like). In binary notation, this mask is shown as follows:

```
11111111.11111111.00000000.00000000
```

For a Class C address, the mask is 255.255.255.0, which means that only the last octet (256 addresses) is available for hosts.



CAUTION: *Host numbers which are all zeroes or all ones are not allowed, hence a Class C address cannot be xxx.xxx.xxx.0 or xxx.xxx.xxx.255, giving 254 usable addresses.*

The various notations for IP addresses and their masks are illustrated in [Figure 2-1](#).

		Network	Host		
Class A 10.15.16.11		00001010	. 00001111	. 00010000	. 00001011
		10	15	16	11
Mask		255	0	0	0
		11111111	. 00000000	. 00000000	. 00000000
		Network	Host		
Class B 129.10.2.3		10000001	. 00001010	. 00000010	. 00000011
		129	10	2	3
Mask		255	255	0	0
		11111111	. 11111111	. 00000000	. 00000000
		Network	Host		
Class C 202.15.23.11		11001010	. 00001111	. 00010111	. 00001011
		202	15	23	11
Mask		255	255	255	0
		11111111	. 11111111	. 11111111	. 00000000

Figure 2-1 IP Address Notation

Valid network numbers for Classes A to D are given below where *hhh* represents the host portion of the address which is assigned by the network administrator.

- Class A: 001.*hhh.hhh.hhh* to 126.*hhh.hhh.hhh*
- Class B: 128.001.*hhh.hhh* to 191.254.*hhh.hhh*
- Class C: 192.000.001.*hhh* to 223.255.254.*hhh*
- Class D: 224.000.000.000 to 239.255.255.255

Subnetting

IP addresses consist of a 32-bit address field, which is divided into two parts: the network identifier and the host identifier. This addressing scheme creates a two-level hierarchy with two major benefits:

- Routing tables only have to contain routes to each network (not to each host).
- Host addresses can be assigned by a local administrator (not a central site).

However, the increasing popularity of the TCP/IP protocol suite and the explosive growth of the Internet created problems with this two-level addressing hierarchy.

- Local administrators had to request a new network number from the Internet when a new network was installed.
- There was tremendous growth in the size of routing tables maintained by IP routers.

These problems were solved by adding a further level of hierarchy to the IP addressing structure. Instead of a two-level (network, host) hierarchy, a three-level (network, subnet, host) hierarchy was created. Each organization is now assigned one (at most a few) network number from the Internet. The organization is then free to assign a distinct subnetwork number for each of its internal networks. This solves the first problem of required registration of network numbers for new segments.

The second problem is solved by guaranteeing that the subnet structure of a network is never visible outside the group of networks implementing it. The route from the Internet to any subnet of a given IP address is the same, no matter which subnet the destination host is on. The local routers need to differentiate between subnets, but as far as the IP routers outside of the autonomous system (AS) are concerned, all of the subnets in an autonomous system are collected into a single routing table entry.



Different routing protocols operating at layers three and four may still cause the second problem to occur. If this is the case, specific router configuration is required.

Subnet Masks

A subnet mask allows the host portion of an IP address to be further divided into two parts: the subnet number and the host on the subnet. Basically, masks determine how much of the address relates to the network, and how much relates to the host.

A 32-bit subnet mask defines the division between subnet number and host number (see [Figure 2-2](#)).

IP address	Network Number		Host Number	
Subnet address	Network Number		Subnet Number	Subnet Host Number
Subnet mask	11111111 11111111		11111111	00000000

Figure 2-2 Subnet Masking

- Subnet mask bits which are set to zero (0) identify the subnet host number. Host bits always begin with the *least* significant bit and work towards the *most* significant bit.
- Subnet mask bits which are set to one (1) identify either the original network number, or part of the subnet number. Network bits always begin with the *most* significant bit and work towards the *least* significant bit.

The subnet mask consists of a similar field structure to that of the IP address. For example, a subnet mask of 255.255.0.0 would mean that the first two three-digit bytes of the IP address (the fields masked by 255.255) are to be recognised and used as the network address, and the last two bytes (those set to 0.0) are to be used to identify the host address.

An alternative way of expressing a subnet mask is a single number indicating how many bits of the IP address are to be used for the

network address. For example 255.255.0.0 can be expressed as 16 whilst 255.255.255.192 can be expressed as 24. Some vendors require the use of this notation when configuring bridges and routers.

Worked Example

You have been allocated a Class C IP address of 193.1.2.0. You need to establish two subnets, each of which must support up to 62 hosts.



Remember that with a Class C address, only the last octet is available for hosts.

Define the Subnet Mask

- 1 Express the IP address in binary format:

193.1.2.0 = 11000001.00000001.00000010.00000000

- 2 You need two subnets, so you need two binary digits.



You need two bits (four possible combinations) for two subnets because 00 and 11 cannot be used. Thus, only 01 and 10 are available.

Select the two most significant bits of the *host* portion of the IP address to define the subnets.

11000001.00000001.00000010.**00**000000

- 4 Define the subnet mask with all network and future subnet bits set to one, and all future host bits set to zero.

Network Number: 11000001.00000001.00000010.**00**000000
193.1.2.0

Subnet Mask: 11111111.11111111.11111111.**11**000000
255.255.255.192



This subnet mask must be configured on each host and defined for each router.

Assign the Host Address

You can now identify the range of addresses which can be assigned to hosts on each subnet.

Subnet 1

193.1.2.64 11000001.00000001.00000010.01000000

Low address:

193.1.2.65 11000001.00000001.00000010.01000001

High address:

193.1.2.126 11000001.00000001.00000010.01111110

.

Subnet 2

193.1.2.128 11000001.00000001.00000010.10000000

Low address:

193.1.2.129 11000001.00000001.00000010.10000001

High address:

193.1.2.190 11000001.00000001.00000010.10111110





NETWARE PROTOCOLS

This chapter explains the basics of Novell NetWare® protocols.

NetWare Internetworking Protocols

Novell commands a large share of the networking market and its Internet Packet Exchange (IPX) protocol is also a network layer standard. Like TCP/IP, IPX is a connectionless datagram protocol. Where TCP/IP refers to networked devices as 'hosts', IPX refers to them as 'nodes'.

The IPX Protocol

IPX is Novell's original network layer protocol. As such, it addresses and routes packets from one entreated device to another on an IPX internetwork.

IPX Addressing

IPX has its own system of internetwork and node addressing. For node addressing, IPX uses the physical address assigned to the specific network interface board within the networked device.

The IPX network address is made up of three components:

- a network number
- a node number
- a socket number

Network Number — Each network segment is assigned a unique network number. This number is used to route packets to their destination network. The network number is a 4-byte hexadecimal address and can contain up to eight alphanumeric characters.

Node Number — The node number identifies the device (node) on the network segment. It is used for local packet transmission. This number is identical to the physical address assigned to the interface board that connects the device to the network.

Socket Number — The socket number physically directs the packet to a particular process within the device or node. This process is the ultimate destination of the packet. Each process which communicates on a network has a socket number assigned to it. This socket number provides a quick way of routing within the node.

Novell Routing Information Protocol (NRIP)

Novell IPX uses NRIP (its own version of RIP called Novell RIP or NRIP) for routing purposes. Although it is similarly named to the IP equivalent, it uses a different protocol. NRIP broadcasts datagrams out onto the network every 60 seconds. Upon receipt of an NRIP datagram, a router adds one to the hop count of each route advertised and broadcasts an NRIP datagram to the other networks, with which it is connected.

The cost of a route in an IPX network is determined by the metric known as *ticks*. In a 'LAN only' environment this is the hop count plus one, for example, three hops or four ticks. For an internetwork connected via a WAN or ISDN links, the tick count is normally based on the speed of the WAN link automatically by the IPX routers.

It should be noted that NetWare 3.X and later versions use the concept of 'internal' IPX addresses, which is somewhat similar to network addressing. The internal address refers to the internal network within that server allowing internal processes to communicate. These numbers must be unique for all servers right across the network. Although network servers may appear wired correctly, and in other respects seem to be working correctly, duplicated internal IPX addresses will not allow correct operation.

NetWare has a hop count limitation imposed by the NRIP. On an IPX network a data packet can cross a maximum of 15 routers before being discarded.

Service Advertisement Protocol (SAP)

Novell also added the Service Advertisement Protocol (SAP) to its IPX protocol family. SAP allows nodes that provide services (such as file servers and print servers) to advertise their addresses and the services they provide. SAPs are broadcast from servers every 60 seconds and routers and servers are obliged to listen to SAP broadcast information, store it in their SAP table, propagate it, and respond to workstation requests.



NRIP and SAP broadcasts occur every 60 seconds interleaved by 30 seconds. If using ISDN, it is important to ensure that spoofing is enabled to minimize call charges. See ["IPX/SPX 'Keepalive' Proxy"](#) on [page 7-10](#) for further details.





APPLETALK

Introduction

AppleTalk[®] is fundamentally different from TCP/IP because it is theoretically 'plug-and-play'. There is no need to configure network addresses when connecting new devices to an AppleTalk network. However, in order to be plug-and-play, AppleTalk routers and associated devices generate significant network traffic to keep each other informed. AppleTalk is designed for use in networks where the devices are permanently connected. Consequently, it is usually impractical to drop ISDN line connections once AppleTalk devices are connected, and high ISDN call charges are usually the result.

Some routers use proprietary features to reduce these line costs. These features include spoofing algorithms to allow line connections to be dropped for significant periods during times when no real user data is present. This minimizes the call charges incurred.

About AppleTalk

AppleTalk is a networking system protocol available on all Apple Macintosh[®] ('Mac') computers and a variety of printer hardware. It is also available on other platforms (for example, UNIX[®] and Windows NT[®]) using various third party shareware and commercial packages. The AppleTalk protocol suite encompasses high level file-sharing using AppleShare[®], LaserWriter[®] printing services, and print spoolers, in conjunction with lower level data streams and simple datagram delivery.

The term 'AppleTalk' was originally used for both the protocol and connecting cables. When the protocol was introduced on different media, the simple shielded twisted pair cable used to connect Macs to other Macs or printers was named LocalTalk[®]. AppleTalk via an

Ethernet is known as EtherTalk[®] and AppleTalk via a Token-Ring network is known as TokenTalk[®]. AppleTalk data can also be carried within other protocols, such as IP via the Internet. 'Encapsulation' or 'tunneling' methods can be used.

AppleTalk Addressing

Devices on AppleTalk networks are known as entities. Each entity on the network has an AppleTalk address consisting of a node number and a network number.

Network Nodes

Node numbers can range from 1 to 253 (254 on LocalTalk) and occupy a single byte. Network numbers are 2 bytes long and can range from 1 to 65535. This can also be written in dotted decimal notation as 0.1 to 255.255 and is similar to writing a 4-byte IP address in the form 128.250.1.21.

Network Numbers

On EtherTalk and TokenTalk extended networks, a network range may be assigned to the cable. This means that AppleTalk nodes on the cable are free to choose a network number from within the specified range.

Network ranges may be zero width (0.5 - 0.5), or larger (0.5 - 0.8). Care should be taken to choose a range with sufficient room for future expansion, but without wasting address space that may be necessary when a connection to another AppleTalk internet becomes available.

Theoretically, the maximum number of AppleTalk nodes that can be accommodated on a single extended network cable is 16,580,355 (65,535 x 253). LocalTalk networks are non-extended and may be assigned only a single network number. The theoretical upper limit for the number of nodes on a LocalTalk network is 254.

In reality, there are physical limitations on the length of each type of cable and the number of possible electrical connections to it.

Zone Names

For convenience, network numbers can be grouped together and described by a 'zone' name. Look-ups for AppleTalk entities in a specific zone generate a lot of traffic, but only on cables that contain those network numbers. A routed network may have up to 255 zones assigned to a single cable. One of these must be denoted as the default zone.

A Mac on a multiple-zone extended network can reside within any one of the available zones (selected by using the AppleTalk Control Panel with Open Transport, or the Network Control Panel on a Mac with classic AppleTalk networking). Non-routed networks are limited to a single zone name per cable.

Within a single physical node, different programs or services can open AppleTalk sockets. The full network, node, and socket address is necessary to specify completely the final destination of an AppleTalk packet.

AppleTalk Named Objects

To assist in finding and distinguishing between different AppleTalk services, an AppleTalk address can be associated with a descriptive name using Name Binding Protocol (NBP). Each entity or service can register an NBP object name and an object type within a zone. For example, a laser printer might register as:

MainRIP:LaserWriter@THE-Printers

Where:

MainRIP is the object name

LaserWriter is the object type

THE-Printers is the Zone Name

Each of the object, type and zone fields are limited to 32 characters in length.

An AppleTalk address for MainRIP might be 73.194/250/129. This shows that the printer process running on node 250 on network 73.194 is listening for printing requests on socket number 129.

A Mac user usually only encounters object and zone names. The Chooser takes care of looking up NBP types and mapping the results to AppleTalk addresses.

AppleTalk Routers

An AppleTalk Router allows AppleTalk services visible on one network interface, such as the built-in Ethernet port, to be used by other hosts connected to a different interface, perhaps on a LocalTalk cable plugged into the router printer port.

Maintaining zone names, looking up NBP names within zones, propagating network routing information, and sending packets between different network interfaces is the responsibility of one or more AppleTalk routers.

Seed and Non-Seed Mode

When multiple routers are connected to the same network, they may all be configured with the same network range and zone name information (they are all seed routers). Alternatively, a router may start up in non-seed mode and obtain configuration information from other routers that are already operating. Once running, there is no practical difference between a seed and a non-seed router.

It is particularly important that all routers connected to a cable have the same configuration information for each of the network range, default zone and zone lists. If this is not the case, then the network may be unpredictable; certain nodes may not be visible or connections may be lost. Many routers handle this potentially serious situation by refusing to start up.

A non-seed AppleTalk entity that starts-up on an extended network initially uses the network number start-up range of 255.0 to 255.254 (65280 to 65534). This network range is used until a router is contacted and the real network range is determined. Network ranges on different physical cables must not overlap, and therefore routers should not be configured with network numbers in the range 255.0 to 255.254.



BRIDGING INTERNETWORKS

This chapter explains the concepts and practical implications of bridging internetworks.

Introduction

Bridges and routers were first used to extend the area a network could cover by connecting two adjacent LANs. Both Ethernet and Token Ring LAN topologies specify limits on the maximum distances between devices and a maximum number of stations that can be connected to a single LAN environment. This distance may be increased with the addition of a bridge or router.

More recently, bridges and routers have been used to segment LANs for performance reasons. When users on a single LAN begin to experience slower response times, the reason is often too much traffic on the LAN. One way of dealing with this is to split a large LAN with many users into several smaller LAN segments, each with fewer users. This increases the routing capacity available to the end user.

Bridges and routers are both devices used to link different LANs or LAN segments together. Many organizations have LANs located at sites that are geographically distant from each other. By placing a router or bridge on the LANs at two distant sites and connecting them with a telecommunications link, users on one LAN can access resources on the other LAN as if those resources were local.

Bridging Versus Routing

Bridging is often looked upon as the poor relation to routing. However, routing and bridging accomplish a similar task in different ways. The primary difference between the two is that bridging occurs at layer two (the datalink layer) and routing occurs at layer three (the

network layer) of the OSI reference model (see [page 1-1](#)). This means that routing and bridging use different information to move packets from one place to another. Bridges and routers both forward packets of information, but routers also determine the path that these packets take. In practical terms, bridging is generally quicker and simpler to configure than routing. However, routing offers greater resilience and control. For more detailed information on routing, see [Chapter 6, “Routing Internetworks”](#).

Bridging Concepts

The simplest bridged network connects two or more LANs (see [Figure 5-1](#)). The interface between the bridge and each LAN segment is known as a port. Each LAN attached to a port is called a network segment.




Figure 5-1 A Simple Bridged Network

The ‘local’ bridge examines each incoming frame and makes a forwarding decision based on the information it contains. When the frame destination is a device on a network segment other than the one on which it was transmitted, the bridge forwards the frame to that port. By forwarding only frames addressed to devices on other segments, bridges increase the effective throughput of the network.

Bridging and the OSI Reference Model

Bridging occurs at layer two of the OSI reference model—the link layer. This means that bridges see the network merely as a collection

of source and destination addresses. They have no knowledge of the paths between addresses and they do not examine any upper-layer information. This means that they can interconnect incompatible higher level protocols such as TCP/IP and DECnet. This does not mean that a DECnet network can receive a TCP/IP encoded frame, but it can forward it to its destination. Bridges can rapidly forward traffic representing almost any network layer protocol.

A bridged network has several benefits.

- More devices can communicate over a bridged network than would be supported on any single LAN connected to the bridge.
- Bridges extend the effective length of a LAN, allowing remote sites to be connected.
- The amount of data forwarded by the bridge is kept to a minimum so devices do not receive large amounts of irrelevant data.
- To some extent, the bridge acts as a firewall for network errors.

Transparent Bridging

Transparent bridges are designed to enable frames to move back and forth between two network segments running the same MAC layer protocols. This type of bridging is called transparent because the end stations are not aware of the existence of intermediate bridges. Transparent bridges have three useful capabilities:

- learning
- filtering
- forwarding

Learning

A bridge's learning capability allows it to prevent unnecessary traffic from flooding the network. When a bridge is first powered on, it does not know the number or addresses of the devices on the LAN connected to it. In order to function correctly, the bridge must learn this network topology. It does this by examining all incoming frames, and building an address table of all the devices it knows to be on the

various segments of the local LAN. All basic bridge functionality involves transactions using this address table (often called the Filter Address Table).

A bridge examines the source address of each frame it receives. It compares this address to the entries already in the source address table. If the address is not there, the bridge adds it. Using this method, the bridge learns the addresses of all the devices on the network. This learning capability allows new devices to be added to the network without reconfiguring the bridge.

Filtering

A bridge allows users to reach any part of the network that they need to, but to minimize traffic on the network, it must be interconnected in such a way that frames are 'filtered' and only those frames that need to pass from one LAN to the other are forwarded across the bridge. Typically, about 80 percent of the frames transmitted on a typical workgroup or department LAN are destined for stations on the local LAN.

Bridges make a simple forward/don't forward decision on each frame they receive from the LAN. This decision is based on the destination address of the frame. If a frame's destination address is on the same LAN segment as its originating address, it is filtered out and not forwarded across the bridge. If it is destined for an address on another LAN segment, it is forwarded over the bridge. Bridges can filter and forward frames very quickly, making them good for large traffic volumes.

Bridges can filter frames based on any link layer field. For example, a bridge can be configured to reject all frames from a particular network. Unnecessary broadcast and multicast frames can also be filtered in this way. Data-link information often includes a reference to an upper layer protocol, and bridges can usually filter on this parameter too.

Forwarding

Forwarding ensures that a frame takes the correct next step to get where it is going. If the destination address is on a different network segment, the bridge determines which of its ports is associated with that address and forwards the frame to the appropriate port. If the destination address is not in the address table, the bridge forwards the frame to all its ports except the one on which it was received.

Transparent bridges are not allowed to forward frames containing errors. They must verify checksums and if an error is detected, the frame is discarded.

Active Loops

Learning, filtering, and forwarding functions rely on the existence of a single path between any two devices on the network. In simple topologies it is relatively easy to ensure that only one path exists. As the number of connections increases or the network becomes more complex, the probability of inadvertently creating multiple paths or 'active loops' between devices increases dramatically.

Active loops can be a severe problem for bridge-based networks. The loops can lead to unnecessary duplication of frames and this redundant traffic can quickly degrade overall network performance. [Figure 5-2](#) illustrates a topology containing active loops. Every time Host A sends a frame to Host C a separate instance of the frame is forwarded by each bridge, resulting in two identical frames traversing the network.

Broadcast Storms

Broadcast frames are delivered to all devices on the network. They are used by Network Operating Systems to advertise file and print services to clients (for example, IPX's Service Advertising Protocol). A *broadcast storm* is a burst of this broadcast frame traffic.

Referring to [Figure 5-2](#), assume that Host A's initial frame is a broadcast. Because there is an active loop on the bridged network, both bridges forward the frame endlessly. This uses all available

network bandwidth and blocks the transmission of other packets on both segments.




Figure 5-2 Active Loops in a Bridged Network

Incorrect Learning of MAC Addresses

Using [Figure 5-2](#) again, assume Host A sends a frame to Host C. Both bridges receive the frame on their Network 1 interfaces and ‘learn’ that Host A is on Network 1. However, when Host C receives two copies of the frame (one from each bridge), both bridges receive the frame again, this time on their Network 2 interfaces, because all hosts receive all messages on broadcast LANs.

The bridges may relearn the address of Host A as being on Network 2. If this is the case, when Host C replies to Host A’s frame, both bridges will reject the frame, because their address tables will indicate that the frame’s destination (Host A) is on the same network segment as the frame’s source (Host C).

The problem of active loops can be addressed by using the Spanning Tree Algorithm. This is now a basic part of bridge functionality.

Spanning Tree

The Spanning Tree Algorithm, sometimes referred to as the Spanning Tree Protocol (STP), creates a set of device to device paths through the network, such that there is only one active or 'primary' path between any two devices. All paths not selected by the spanning tree are temporarily disabled. In other words, the Spanning Tree Algorithm (STA) creates a logically loop free network topology by using certain paths and blocking others. In [Figure 5-3](#) the diagram on the left shows a meshed topology containing loops. The figure on the right shows how a spanning tree can be placed over the meshed topology to automatically eliminate these loops.




Figure 5-3 Example of Spanning Tree Algorithm

Spanning tree allows participating bridges to reactivate blocked paths if an existing primary path fails. With this feature, the STA allows networks to recover quickly and automatically if a network device such as a bridge or a section of network cabling fails.

Spanning Tree Problems

Although the STA solves many problems, it can also create them for wide area networks. If there are active loops in the long distance part of the network, the STA will disable one or more lines to eliminate them. However, even though a line is disabled, the physical connection remains intact. Because long distance lines are most often leased, network managers who choose bridging as a basis for wide-area internetworking may find themselves paying for long distance lines that are not actually used because they have been disabled by the STA.



CAUTION: *Spanning Tree is recommended for permanent links only. It should not be used in ISDN networks and other semi-permanent connections.*



Token ring and FDDI networks can also implement Source Route bridging. This is an IBM standard that routes frames by specifying forwarding information in the frames themselves.

Local and Remote Bridging

Bridges may be either local or remote. Local bridges connect multiple LAN segments within the same local area. Local bridges connect to local transmission media, particularly network backbones. Typical media include coaxial, fiber-optic, and twisted pair cable, so a local bridge may have more than one physical LAN port on it.

Remote bridges are also known as wide-area bridges. Remote bridges connect multiple LAN segments in different areas. Remote bridges usually use ISDN or telephone lines to connect these remote LAN segments. Remote bridges often only have one physical LAN port on them, with other ports associated with leased line or dial-up WAN link connections.

They connect to remote access media. There are two basic types of remote access technology. These are discussed in greater detail in [Chapter 7 "Remote Access Using ISDN"](#).

Advantages and Disadvantages of Bridging

Advantages

When deciding whether to implement bridging on a network, consider the following advantages.

- Bridges are simple to install. Advanced bridging features can be implemented with minimum configuration.
- Bridges are transparent to users.
- Bridge-based internetworks adapt automatically to network changes, and can be modified and reconfigured easily.
- Bridges can connect networks running different protocols without requiring additional software. They operate below the network layer in the OSI reference model, so it is not necessary for


network managers to know in advance which high-level protocols will be used.

- Some protocols are 'unroutable', such as DEC's Local Area Transport (DECLAT) protocol which is used for terminal communications. These unroutable protocols must be bridged.
- Bridges form single logical networks. All of the interconnected network segments have the same network identifier. This means that devices can be moved around within the network without configuring new network addresses for them.

Disadvantages

- Bridges cannot load-split over network segments. This means that they cannot take simultaneous advantage of redundant paths in a network.
- Bridges may propagate significant increases in network traffic at certain times and flood the network. For example, this can occur when a frame with an unknown address is sent out.
- Bridges cannot prevent 'broadcast storms'. A broadcast storm may occur when certain broadcast protocols cause frames to be flooded to every port. If there is a malfunction or an incorrectly configured parameter, these activity spikes can be severe enough to render the entire network inoperable.
- Bridges do not provide significant support for fault isolation or other distributed management capabilities. Networks become harder to manage and maintain as their size and complexity increases. Bridges form a single logical network often making fault isolation in very large bridged networks almost impossible.
- Bridge-based internetworks may require extra attention from network administrators to track what is running on the network and where.
- Using bridges to connect networks across wide area fixed links (or leased lines) can cause a problem. If the line speed of the wide area link is too slow, applications on the end stations may timeout causing unnecessary retransmission of frames. If this situation is likely to occur, routers should be used for the remote link.





ROUTING INTERNETWORKS

This chapter explains the concepts and practical implications of routing internetworks.

Introduction

Like bridges, routers consolidate two or more networks into an internetwork. Unlike bridges however, routers maintain the logical identity of each network segment. Therefore, an internetwork based on routing consists of many different logical subnetworks, each of which is a potentially independent administrative domain.

Routers are more complex than bridges. They use the Network Layer Protocol information within each packet to route it from one LAN to another. The router must be able to recognize all of the different Network layer protocols which may be used on the networks it is linking together. This is where the term 'multiprotocol router' comes from—a device that can route using many different protocols. The most common multiprotocol routers route IP and IPX. Routers share information with each other allowing them to determine the best route through a network that links many LANs.

Routing Concepts

Whilst a bridge examines all frames sent on its attached network segments, a router receives only packets specifically addressed to it. This means that routers have more decisions to make than bridges, and they need more information with which to make them. This additional information is contained in the routers routing tables.

A router has two basic functions.

- It must create and maintain the routing tables.

- It must select the next leg of the journey for each packet it processes. This path is selected based on the information contained in the packet and in the routing table appropriate to that packet.

Routing Tables

A routing table contains a variety of information including destination/next hop associations and path desirability. Next hop associations tell a router that a particular destination can best be reached by sending the packet to a specific router which represents the 'next hop' on the way to the final destination. When a router receives a packet, it examines the destination address and associates it with an appropriate next hop.

Path desirability concerns the most efficient path a packet can take. The source and destination routers compare routing metrics to determine the most desirable path between them.

A routing metric is a standard of measurement used by routing algorithms to determine the most efficient path to a particular destination. Routing algorithms store route information in routing tables. This information varies with the routing algorithm used.

Routers communicate with each other by sending messages. These messages include routine updates to routing tables. By analyzing these updates, each router on the network learns the network topology and updates its own routing tables accordingly. This process occurs automatically. The Routing Information Protocol (RIP) is most commonly used for this

Static Routes

The network administrator can define 'static' routes if necessary (for example, if a particular routing policy needs to be enforced). Static routes force traffic through the network in a specific way.

The disadvantage with static routing is that if the network links in the routing definition are down, traffic cannot be passed. The implementation of a static route prohibits the router from offering an alternative data path.

Switching

Switching algorithms are similar for most routing protocols. A host determines that it must send a packet to another host. The source host sends a packet to a router's physical (MAC) address, but with the protocol (network) address of the destination host.

The router examines the packet's destination protocol address to determine whether it knows how to forward the packet to the next hop. If the router knows how to forward the packet, it changes the destination physical address to that of the next hop and forwards the packet. If the router doesn't know how to forward the packet, it drops it.

The packet is forwarded in this way until it reaches its final destination. Although the packet's physical address may change many times, its protocol address remains the same.

Routing and the OSI Reference Model

Routing generally occurs at layer three of the OSI reference model — the network layer. It involves two basic activities: the determination of the best path (routing) and the transport of information over the network (switching). Switching is relatively straightforward, but determining a path can be complicated.

Routers do not connect at the data link layer of the OSI reference model, so they can connect network environments which have dissimilar addressing structures (assuming they have interfaces to each LAN type).

Routers are visible to end stations. This allows them to control the flow of traffic from a transmitting station to a receiving station. If the transmitting station sends packets faster than the receiving station can store them in its buffer, some routers may also be able to signal the transmitting station to stop or slow the transmission, thus controlling the flow and avoiding congestion.

Bridge/Router

Many vendors have created devices that mix bridging and routing technology together in a single system. A bridge/router can act as both a bridge and a router at the same time. Few pure routers are available, because there will always be a need to bridge unroutable protocols.

Figure 6-1 illustrates how a multiprotocol bridge/router processes the packets that it receives.




Figure 6-1 Multiprotocol Bridge/Router

TCP/IP traffic is sent to the IP routing module for processing. Routing is based on the destination IP address contained in each packet.

IPX traffic is sent to the IPX routing module for processing. Routing is based on the destination network number contained in each packet.

All other traffic (frames) is sent to the bridge module for processing. Forwarding is based on the destination MAC layer address contained in the frame. The bridge module does not examine the network protocol address.

Routing Protocols

Routers communicate with each other through protocols that operate at the network layer level. These routing protocols determine whether routing tables are static or dynamic, whether link-state or distance-vector routing is used, and other variables that pertain to communication between routers.

Most routers are dynamic; building and maintaining their routing tables automatically. Often, the facility for a network administrator to manually add in defined static routes is also available.

IPX Routing

Routing on a Novell network is through the Internet Packet eXchange (IPX) protocol suite. The routing protocol is called Novell Routing Information Protocol (NRIP). NRIP-IPX uses distance-vector routing and the maximum number of hops is configured on each router with the default set at 15. Each end station must send an NRIP request to determine which router is best for a desired network. Novell distinguishes between internal routers (those that exist as software on a NetWare server) and external routers (those that are standalone dedicated routers).

Novell also offers a routing protocol called NetWare Link Services Protocol (NLSP) which uses link-state routing. NLSP addresses some problems caused when using NRIP-IPX. For example, the volume of overhead generated from NLSP is much less than from NRIP-IPX.

IP Routing

The TCP/IP protocol suite contains a large number of standards-based protocols. Routing protocols include Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Exterior Gateway Protocol (EGP) and Border Gateway Protocol (BGP). Although IP does not use distance-vector routing, it does impose a 15 hop restriction on routes.

Advantages

- Routers can eliminate traffic on a network because they do not forward broadcast packets from one segment to another.
- Routers are generally more flexible than bridges. They can differentiate between different paths based on factors such as cost, line speed, and line delay, and can be configured (for example) for equal-cost load splitting. Router-based networks may be customized to more closely reflect business requirements.
- Routers can provide a protective firewall between subnetworks. This prevents incidents that occur within one subnet from affecting others, and makes large routed networks easier to maintain than their bridge-based equivalents.

- Router-based networks support any topology, and can more easily accommodate greater network size and complexity than similar bridged environments.
- Routers provide and can take advantage of redundant network paths, allowing them to load split certain applications, helping to ensure that available bandwidth is optimally exploited. Bridges cannot normally do this because the spanning tree algorithm has to be applied and it blocks redundant paths.
- Some routers can translate packets from one data link layer protocol (such as Ethernet) to another (such as Token Ring) much easier than bridges can. These are sometimes called 'gateway routers'.

Disadvantages

- Routers are protocol dependent devices, so they require software for each protocol they run.
- The more protocols a router supports, the more knowledge network administrators must have to configure and troubleshoot the router. Personnel with adequate training may not always be available.
- If it is running a static protocol, configuring a router can be a laborious, time consuming process.
- Some protocols that operate below the network layer are not routable and must be bridged.
- Troubleshooting and diagnostics on a routed network requires a much higher level of expertise than that needed for an equivalent bridged network.



REMOTE ACCESS USING ISDN

This chapter provides an introduction to the basic concepts of ISDN. It describes and illustrates how ISDN can be used for remote access purposes.

Introduction

Integrating voice and data networks can reduce costs and expand capabilities. Integrated access to voice, video, and data services also provides access for applications such as desktop videoconferencing.

ISDN (Integrated Services Digital Network) offers many benefits for organizations where data applications use public switched telephone network facilities. These benefits make ISDN particularly attractive for small regional and international branch sites that need to connect to central enterprise networks and to one another.

- ISDN can carry multiple services—voice, video, and data—on a single network over existing twisted-pair copper wire, so telecommunications service providers and subscribers can dramatically reduce their infrastructure and maintenance costs.
- Basic Rate Interface (BRI) ISDN provides much higher bandwidth than analog modem-based solutions. Using ISDN compression ratios from 2:1 to 4:1, it is possible to deliver effective transmission rates ranging from 256Kbps to 632Kbps.
- ISDN provides a clearer, less noisy voice telephone service, with the built-in security advantage of digital transmission and managed by easy-to-use call control features (dependent on the access devices used).
- Remote users working from home or on the road can use high-speed ISDN to access critical central site resources at higher performance levels.

- ISDN caller identification features can enable some ISDN access devices to screen incoming calls based on the caller's phone number ID, and accept or reject the call based on user-specified preferences. It can link to a directory and forward the call accordingly, or map to a database to pull the caller's record. It can even bypass the local site and link the call to a remote-site IP address for routing purposes.
- A dynamic bandwidth allocation feature included in some ISDN access devices can aggregate data channels in real time to accommodate even the most bandwidth-intensive applications.
- An ISDN connection can act as a low-cost backup for a leased line on a "pay only for use" dial-up connection basis with line speeds comparable to current T1/E1 leased lines (where ISDN Primary rate Interface services are employed with up to 30 B-channels). ISDN eliminates the expense of a second leased line that may go unused.
- ISDN can handle multiple devices on a single line. Up to eight telephones, computers, workstations, faxes, credit card readers, cash registers, or other devices that are connected via an ISDN access device (or are ISDN compatible in themselves) can be directly attached to a single ISDN line.



Although several devices can be supported on a single line, only two simultaneous connections can be supported on BRI ISDN.

- ISDN's end-to-end digital transmission delivers more accurate and reliable connectivity than analog technology. It normally has lower error rates and fewer dropped connections.
- ISDN's technology provides quicker connect times to better support LAN protocols such as IP and IPX, which require lower latency across the connection. This is particularly useful for Internet and other on-line services as well as in retail credit verification applications.
- ISDN interoperates with other WAN services such as existing analog services, X.25, Frame Relay, Switched Multimegabit Data Service (SMDS), and higher-speed services like ATM.

- ISDN can provide dial-up access for IBM users: for example, cluster controller to front-end processor (FEP) links and inter-FEP links at T1 channel extension rates.
- Attractive tariffs and expanded availability make ISDN a cost-effective alternative to private leased lines for low- and high-speed data networking.

How ISDN Works

In an analog network, a two-wire loop from the telephone company's local exchange to the customer premises supports a single transmission channel which can carry only one service (voice, data, or video) at a time. With ISDN, this same pair of twisted copper wires is logically divided into multiple channels.

ISDN Logical Channels

ISDN defines two types of logical channels. They are distinguished by both function and capacity:

- B (bearer) channels operate at 64 Kbps and carry circuit-mode or packet-mode user information such as voice, data, fax, and user-multiplexed information streams. All network services are available through B channels.
- The D (data) channel operates at 16 Kbps for BRI and 64 Kbps for PRI. It carries call signaling and setup information to establish a network connection, request network services, route data over B channels, and close the call when complete. This information is designed to travel through a totally separate, dedicated communications network from the bearer channels. It is this out-of-band signaling network that gives ISDN faster connection times—from one to four seconds as opposed to 10 to 40 seconds for analog dial-up lines. On some ISDN networks (where the service provider allows) bandwidth not required for signaling and control on the D channel can be used to transport user packet or frame data when needed.




Figure 7-1 BRI and PRI Interfaces

ISDN User Interface Standards

Users connect to ISDN by means of a local interface to a "digital pipe." ISDN supports digital pipes of various sizes to satisfy different application needs. For example, a residential user might require enough capacity to handle a telephone and a PC. However, a remote site connecting to ISDN via an on-premises private branch exchange (PBX) or a bridge/router might require a higher-capacity pipe. At different times, the pipe might use varying numbers of channels, up to its capacity limit.

The ITU-TSS has defined two ISDN user interface standards:

- The Basic Rate Interface consists of two B channels and one D channel for signaling (2B+D).
- The Primary Rate Interface is a 30B+D interface (23B+D interface in North America and Japan). It is the ISDN equivalent of the 2.048 Mbps (or 1.544 Mbps) interface over a T1/E1 line or trunk; the physical layer is identical for both. The D channel is channel 16 (or 24) of the interface, and it controls the signaling procedures for some or all of the B channels.

ISDN for Remote Access

Interconnected LANs may be in the same geographic area or they may be separated by great distances. When they are geographically distant, they are connected into a Wide Area Network (WAN). There are numerous methods of creating a WAN for remote access, but they all fall into one of two categories: dedicated or switched.

Dedicated access — provides a constant data transmission path between two *specific* points. This is typically in the form of a leased or private line. The line is always available, and large amounts of data can be sent 24 hours a day if necessary. However, the cost is high—the connection costs the same whether it is used 24 hours a day or just one hour a day. In practice, many costly leased lines are only used for a short time each day.

Switched access — provides a connection between *any* two points on an as-needed basis. An ISDN line or a standard analog telephone line can be used to provide switched access. Anywhere can be accessed as and when needed. Switched access is a low cost solution for intermittent users—the cost is normally based on the line time used plus a fixed rental charge. However, the cost will often exceed that of a dedicated leased line if the connection is used for 24 hours per day.

ISDN can be used as the sole means of remote access or as a backup if the permanent leased line has failed.

Primary Link Backup

In this application, the bridge/router normally communicates with the central site via a dedicated fixed link such as a leased line. If the fixed link fails, the router automatically uses the ISDN line to dial up the central site and resume communications. When the fixed link is restored, the ISDN link is automatically terminated. See [Figure 7-2](#).




Figure 7-2 ISDN as Backup to a Leased Line

Dial on Congestion

This application is similar to the fixed link backup. The bridge/router normally communicates with the central site via a dedicated fixed link such as a leased line. When network traffic exceeds the available bandwidth, the bridge/router automatically uses the ISDN line to provide additional bandwidth ('bandwidth top-up'). When traffic levels fall below a programmed level for a set time, the bridge/router automatically terminates the ISDN call and all traffic returns to the fixed link. See [Figure 7-2](#).

Dial-on-Demand

In this application the bridge/router has no dedicated connection to the central site. The bridge/router only calls the central site via ISDN when there is data to be transmitted. When the connection has been established, it can be held open for other traffic. If there is no traffic for a programmable amount of time, the bridge/router drops the call. See [Figure 7-3](#)

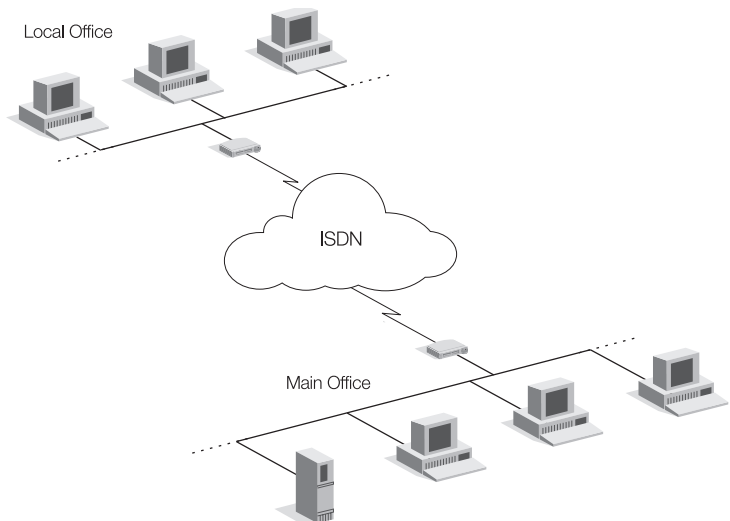


Figure 7-3 Dial-on-Demand

Telecommuting

Working at home provides many benefits to both the employee and the employer. The benefits of telecommuting are possible using widely available applications that BRI ISDN can support. The applications include remote LAN-access, file transfer, terminal emulation, and screen sharing applications. The ISDN line can also support simultaneous voice applications.

In [Figure 7-4](#), the telecommuter connects both a PC and an analog phone to the bridge/router. The bridge/router uses both B channels for data unless there is an incoming or outgoing voice call. When a voice call is in progress, the bridge/router uses a single B channel and the voice call uses the other one. This enables the telecommuter to have a single incoming voice/data line.




Figure 7-4 Telecommuting

Security Management

Implementing ISDN access to a private network opens that network to potential unauthorized access. To minimize this threat, some ISDN access devices allow a security procedure to be implemented using the PPP link protocol. Two such procedures are Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP). Needless to say, these security features do add an overhead in the data volume when operating and extra calls are established.

PAP — PAP provides one way authentication of the remote bridge/router calling into a central site. The remote bridge/router gives a user name and password to be validated by the central site bridge/router.

CHAP — CHAP uses three way handshake authentication. One side of the connection initiates the challenge with a key to a 'secret'. This secret is used to determine a response which is encoded and sent back to the initiator. The response is evaluated and either accepted or rejected. Additional challenges may be issued throughout the session to confirm that the caller is valid. The nature of the challenge is random to avoid duplication of response.

Minimizing Costs

ISDN lines are normally charged by connection time (plus a fixed rental). However, the LAN protocols used with them were developed for the local environment where bandwidth is essentially free. Bridge/routers with ISDN interfaces must not only connect to the ISDN network to route the data, but must also optimize the use of that network to minimize connection charges.

Data Filters — Although there may be many devices on a remote network, some of these devices may never need to connect to any other location. Filters in some ISDN bridge/routers can be used to allow only authorized users to contact a remote connection.

Data Compression — When a remote connection has been established, it is practical to maximize the use of the available

bandwidth. A bridge/router with data compression can transport up to four times more data across a 64Kbps ISDN link than a basic ISDN terminal adapter. Data compression does not always speed the transmission of data. It is therefore important to have a bridge/router which allows data compression to be switched on and off.

IPX/SPX 'Keepalive' Proxy — Novell is a very 'chatty' protocol with the file server sending a message to each remote client terminal every five minutes when there is no actual user communication. In a network where bridge/routers are using ISDN as the transport, these keepalive messages can increase the monthly connection charges significantly. Bridge/routers should implement a 'spoofing' protocol whereby the bridge/router responds to the keepalive messages sent from the file servers without actually bringing up the ISDN connection.

Demand RIP for IP/RIP and SAP for IPX — Normal routing protocols cause bridge/routers to communicate with their neighbors to determine which paths are available to transmit data. This is done via periodic messages sent throughout the network. Most often, nothing has changed in the network and the messages do not convey any new information. Bridge/routers should implement either static routing or demand RIP/SAP. Demand RIP/SAP only sends routing and service updates when there is a change in topology or the status of a particular service changes. This minimizes ISDN connection charges since change information is only sent across the network when an ISDN connection is up.

Multilink PPP — This feature allows a bridge/router to dial up additional ISDN B channels when there is a large amount of data traffic to be sent across the network. The additional bandwidth shortens the transfer time and improves interactive performance.

Some ISDN bridge/routers are capable of building and collapsing 'scalable pipes' across ISDN. In this case, users are only charged for the bandwidth needed to support their applications when they are actually needed.

Timebands — Timebands allow users to establish certain times at which calls are to be automatically placed to particular destinations.

Typically, calls can be scheduled at different times of the day for different days of the week. The call duration can also be specified.

Tokens — Tokens enforce a maximum call use over a given time period, typically a month, to control ISDN usage and line charges.

Priority Queuing — Packets are generally queued for transmission onto the WAN. Using priority queuing, each session is allocated its own queue and these queues are serviced in a 'round robin' manner. This avoids a bandwidth hungry application (such as an FTP session) from starving other processes of bandwidth.

Queues can be organized by IP address or protocol and flexible prioritization schemes can be used to give certain hosts or applications priority.



GLOSSARY

10Base2

An IEEE standard for using IEEE 802.3 protocol at 10 Mbps over thin Ethernet cable.

10Base-T

An IEEE standard for using IEEE 802.3 protocol at 10 Mbps over unshielded twisted-pair cable (the T stands for twisted pair).

802.3

An IEEE standard for the physical layer that specifies a CSMA/CD protocol. This is the standard protocol used for Ethernet. Refer to *CSMA/CD*.

Address

The unique code assigned to each device or workstation connected to the LAN.

ATM

Asynchronous Transfer Mode; a high-speed switching and multi-plexing technology that uses 53 byte fixed-length cells. ATM is the standard switching technology for B-ISDN.

AUI

Attachment user interface, the interface between the unit and the data terminal equipment, usually in the form of a connecting cable.

B Channel

The bearer (B) channel is a 64 Kbps circuit or packet switched channel used for transporting user information: voice, data, images, and video.

Bandwidth

The capacity of data communications system or channel.

B-ISDN

Broadband ISDN; the second generation of ISDN that provides transmission rates greater than ISDN PRI (>2Mbps).

BONDing

An international standard for aggregating multiple data channels into a single logical connection with bandwidth greater than 64Kbps, It is very popular in video conferencing applications and is sometimes known as inverse multiplexing.

BRI

Basic Rate Interface; the ISDN user interface standard that provides two 64Kbps B channels and one 16Kbps D channel, and has physical and logical access to basic rate ISDN.

Bridge

A device that links two or more local or remote area networks together. A bridge may be used to extend the network or to connect two different network transport technologies together.

Broadcast Storm

An event in which broadcast frames are propagated endlessly through the network. Usually due to poorly configured bridge and router connections.

CCITT

Comité Consultatif International Téléphonique et Télégraphique; now renamed ITU, International Telecommunication Union.

Channel Aggregation

Combines multiple physical channels into one logical channel of greater bandwidth. For example, with BRI connections, it combines the two 64 Kbps channels into a single, logical 128Kbps channel.

CHAP

Challenge Handshake Authentication Protocol; part of the PPP protocol to ensure authentication of the connection between two devices.

Class

Type of IP address; IP addresses fall into three main classes, A, B and C.

Clear Channel

A transmission method that can offer the entire 64 Kbps channel bit rate for data transfer because signalling and control information are handled out-of-band (on a separate channel).

Client

A user whom is making use of a particular system resource or peripheral through a workstation attached to a local or wide area network.

Client/server

A user who is attached to a file server to recover and store files, but the processing of the data or use of an application is carried out on the client machine.

Coaxial cable

A twin-conductor cable used for computer networking, in either a thick or thin form. This cable consists of a centre core wire (stranded or single core) covered by insulation, a second conductor of woven wire, and an external covering of rubber. Thin coaxial cable resembles television cable. Thick coaxial cable has an increased diameter outer bore and is often yellow or orange in color.

CSMA/CD

Carrier Sense Multiple Access with Collision Detection. A refinement of CSMA in which stations are able to detect the interference caused by simultaneous transmissions by two or more stations (collisions) and to retransmit colliding messages in an orderly manner.

D Channel

A control channel carrying signaling information, running at 16 Kbps. The basic rate ISDN 2 service carries two B channels plus one control D channel. Refer to *ISDN* and *B Channel*.

Data

Characters or code either entered by the user or passed between devices that are part of the computer or network.

Data communications

The transfer of data via transceiver equipment by means of data transmission according to a protocol. Refer to *Protocol*.

Downloading

A user initiated transfer of data from a server to the user's own workstation. Also used to classify the transfer of files from one system to another, usually to upgrade or revise system software.

Enterprise Network

A geographically dispersed network under the auspices of one organization.

Ethernet

A 10 Mbps baseband local area network protocol, compatible with IEEE 802.3 standards.

Euro-ISDN

The European implementation of ISDN.

FastConnect

The OfficeConnect Remote's proprietary connection protocol that allows fast connection between units either over the ISDN link or over a permanent leased line WAN link.

Fiber optics

A technology that uses laser light pulses, sent over thin glass fibres, which is able to deliver data at speed up to several gigabits per second.

File server

A computer running a special operating system that allows workstations to access files.

Filter

A configuration that removes types of data frames based on user-entered parameters.

Firewall

A method of preventing unauthorized access to a network or a host on a network. A firewall is usually implemented within a router's software.

Frame Relay

A fast packet-switching WAN technology for interconnecting LANs at high speed. Frame relay defines the interface between the user equipment and the WAN; it does not define internal operation of the network or the interface or protocols used within the WAN itself. The method by which a data packet is constructed to be sent across a network. Usually assembled with header and footer information.

Gateway

Another name for a router on a network.

HDLC

High-level Data Link Control. OSI's bit orientated protocol.

Hop count

The number of routing nodes between a source and destination device on a LAN or WAN.

Host

A device or computer on an IP network to which you can connect.

Hub

A cabling centre in a star topology that either amplifies a signal and transmits it (active hub) or simply passes the signal along (passive hub).

Hyperterminal

A terminal emulation program provided with Microsoft Windows 95.

IEEE

The Institute of Electronic and Electrical Engineers.

IPX

Internetwork Packet Exchange, the default data packet protocol for Novell's NetWare operating system.

IPCP

Internet Protocol Compression Protocol

ISDN

Integrated Services Digital Network; a public switched digital network that provides a wide variety of communications services and integrated access to the network.

ISO

International Standards Organization. Refer to *Open Systems Interconnection*.

Kbps

A measurement of data transmission speed in kilo bits per second.

Keep alives

A message sent by one network device to inform another network device that the virtual circuit between them is still active.

LAN

Local area network, a network that covers a group of local workstations and peripherals that require to share information.

Learn

A bridge learns addresses received at any of its interfaces and adds them to its filter address table.

Leased Line

A dedicated and non-switched circuit, typically supplied by the telephone company, that permanently connects two or more user locations. A leased line may be digital or analog, and point-to-point or multipoint.

MAC

Medium Access Control, a protocol for determining which device has access to the network at any one time.

Mbps

A measurement of data transmission speed in megabits per second.

MAN

Metropolitan area network, a network that covers a city.

MIB

Management information base.

Multilink PPP (MLPPP)

A variant of PPP which addresses the additional features of compression and channel aggregation.

NETBIOS

Network Basic Input/Output System, a standard for supporting network communications that is independent of the underlying network transport type.

NetWare

Novell's Network Operating System (NOS) line.

Network

A method of connecting computers and other devices together with cabling so that they can communicate with each other.

NIC

Network interface card, an expansion card that enables a PC to communicate on a network.

Network layer

The third layer of the OSI reference model. This layer is responsible for controlling message traffic.

NFS

A network file system developed by Sun Microsystems for shared files over a UNIX platform.

Node

An alternative name for a computer or device (such as a printer or modem) that is connected to a network.

NOS

Network operating system.

OSI

Open Systems Interconnection, a body of standards set by the International Standards Organization to define the activities that must occur when computers communicate. There are seven layers, and each contains a specific set of rules to follow at that point in the communication.

PAP

Password Authentication Protocol. Part of the PPP protocol to ensure authentication of the connection between two devices.

Peer-to-peer network

A network which contains workstations which are able to act as both client and client servers.

Piggyback

A way of transmitting routing information over ISDN lines by adding it to valid data frames. This avoids ISDN calls being generated solely for passing routing information.

Physical layer

The first layer of the OSI network layer model. This layer manages the transfer of individual bits of data over wires, or whatever medium that is used to connect workstations and peripherals.

Polling

A method of controlling terminals on a multi-point network where each device is interrogated in turn to determine if the device is ready to receive or transmit data.

POTS

Plain old telephone service; the existing analog telephone lines.

PPP

Point-to-Point Protocol; the protocol for routing between devices made by different manufacturers.

Presentation layer

The sixth layer of the OSI network layer model. This layer controls the formatting and translation of data.

Protocol

A set of rules and procedures that govern the exchange of data between two communicating systems.

PSTN

Public switched telephone network.

Quick Configuration

A set of menu driven forms in the management system that allow you to configure the unit for most types of ISDN connection.

REN

Ringer equivalence number.

RIP

Routing information protocol.

Router

A protocol transparent device that links networks. A router can be used to separate unwanted traffic on either side of the bridge, reduce the traffic, or to provide security from unauthorized users.

SAP

Source Advertisement Protocol.

Segment

A section of an Ethernet network, typically connected by repeater or a bridge to another segment.

SPX

Sequenced Packet Exchange, Novell's guaranteed-delivery version of IPX.

Session

A logical connection between two communicating systems that allows for the transfer of data.

Session layer

The fifth layer of the OSI network layer model. This layer is responsible for the security and administrative tasks of communicating.

SNMP

Simple network management protocol, a software program to allow the remote management of bridge and routing devices.

Static Route

A route you have entered and made permanent rather than a route that the unit has learned by connecting to other routers.

STP

Spanning tree protocol, a protocol which prevents network loops.

Terminal

The Microsoft Windows terminal emulation program.

Terminators

Devices that are used at the ends of a linear bus network segment to reflect the signal back and prevent failure of the segment.

TCP/IP

Transmission Control Protocol/Internet Protocol, a set of communication protocols that support peer-to-peer connectivity functions for both local and wide area networks.

Thick Ethernet

A cabling system for Ethernet connections that uses a heavyweight coaxial cable. Suitable for large networks.

Thin Ethernet

A cabling system for Ethernet connections that uses a lightweight coaxial cable. Suitable for small networks.

Ticks

A measure of the time taken to pass information through a routed network.

Token Ring

A network transport technology in which an electronic token that allows access to the network is passed around stations in the ring.

Topology

The way that a network is physically laid out.

Transport layer

The fourth layer of the OSI network layer model. This is responsible for error checking and correction, and some message flow control.

WAN

Wide area network, a network that covers a wide area and requires special communication devices (bridges and/or routers) to make connection possible. WANs make connections over long distances and need telephone, satellite, or microwave equipment to allow connections to be made.

Workstation

Another name for a computer (or terminal) that is connected to a network.



INDEX

Numbers

10Base2 G-1
10Base-T G-1
802.3 G-1

A

active loops 5-5
address 2-6, G-1
 AppleTalk 4-2
 class 2-5
 hardware 1-6
 internet 2-4
 IPX 3-1
 link layer 1-6
 logical 1-6
 MAC 1-6
 network layer 1-6
 virtual 1-6
address notation 2-6
addressing 1-6
allocation programs 1-4
AppleShare 4-1
AppleTalk 4-1
 named objects 4-3
 overview 4-1
 routers 4-4
application layer 1-2
audience 1

B

B channel 7-3, G-1
bandwidth G-1
B-ISDN G-2
BRI 7-1, 7-3, G-2
bridge 1-8
bridge/router 6-4
bridges 5-1, G-2
 advantages 5-8

 concepts 5-2
 disadvantages 5-9
 filtering 5-4
 forwarding 5-5
 learning 5-4
 loops 5-5
 relationship to OSI reference model 5-3
 transparent 5-3
 versus routers 5-2
bridging loops 5-5
broadcast storm 5-5, G-2

C

class
 address 2-5
conventions
 icons 2
 text 2

D

D channel 7-3, G-3
data link layer 1-3
dotted decimal notation 2-6
dynamic routing protocol 2-2

E

Enterprise Network G-4
enveloping 1-5
ethernet G-4
EtherTalk 4-2

F

file server G-4
file transfer 2-4
firewall G-5
frames 1-5

FTP 2-4

G

gateway 1-7

H

hop count 2-3, G-5

host 1-7

I

icon conventions 2

IETF 2-1

IGP 2-2

International Standards Organization 1-1

Internet protocols 2-1

internetworking

environment 1-1

overview 1-1

internetworking devices 1-6

InterNIC 2-5

IP 2-1

address

subnet 2-8

subnet mask 2-9

address notation 2-6

addressing 2-4

IP routing 6-5

IPX 3-1

IPX routing 6-5

ISDN 7-1, G-6

B channel 7-3

D channel 7-3

dial on congestion 7-6

dial-on-demand 7-7

interface standards 7-4

PRI 7-3

remote access 7-5

telecommuting 7-8

ISO 1-1

L

leased line G-6

logical address 1-6

loops 5-5

M

MAC G-6

MAC address 1-6

incorrect learning 5-6

mask 2-6

subnet 2-9

metrics 2-3

MLPPP G-7

multiple protocols 1-5

N

NetWare protocols 3-1

network layer 1-3

network layer address 1-6

network number 3-1

node number 3-2

NRIP 3-2

O

OSI reference model 1-1

application layer 1-2

data link layer 1-3

network layer 1-3

physical layer 1-3

presentation layer 1-2

relationship to protocols 1-4

session layer 1-3

transport layer 1-3

P

packets 1-5

physical layer 1-3

point-to-point 2-3

POP 2-3

PPP 2-3

presentation layer 1-2

PRI 7-3

protocol 2-3, 2-4

AppleTalk 4-1

dynamic routing 2-2

IPX 3-1

NetWare 3-1

Novell Routing Information 3-2

routers 6-4

Service Advertisement 3-3

protocol (*continued*)
 terminal emulation 2-4
protocol stack 1-5
protocol suite 1-5
protocols 1-4

R

remote access 7-5
repeater 1-8
RFCs 2-1
RIP 2-2
routers 1-8, 6-1, G-9
 advantages 6-5
 concepts 6-1
 disadvantages 6-6
 protocols 6-4
 relationship to OSI model 6-3
routing
 IP 6-5
 IPX 6-5
routing table 6-2

S

SAP 3-3
session layer 1-3
significant bit 2-9
SLIP 2-3
SMTP 2-3
SNMP 2-4
socket number 3-2
spanning tree 5-7
SPX G-9
subnet mask 2-9
subnetting 2-8
switching 6-3

T

TCP 2-2
TCP/IP 2-1, G-10
telecommuting 7-8
telnet 2-4
text conventions 2
ticks G-10
TokenTalk 4-2
transport layer 1-3

U

UDP 2-2

V

videoconferencing 7-1

W

WAN G-11

Z

zones 4-3