INTRODUCTION TO CISCO ROUTER CONFIGURATION: STUDENT GUIDE

CISCO INTERNETWORK OPERATING SYSTEM

Release 11.2

CISCO SYSTEMS

Introduction to Cisco Router Configuration: Student Guide

Cisco Internetwork Operating System Release 11.2

Corporate Headquarters
170 W. Tasman Drive
San Jose, CA 95134-1706
USA
408 526-4000
800 553-NETS

Text Part Number: 78-3774-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

SOFTWARE LICENSE

PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THE SOFTWARE. BY USING THIS SOFTWARE YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS LICENSE. IF YOU DO NOT AGREE WITH THE TERMS OF THIS LICENSE, PROMPTLY RETURN THE UNUSED SOFTWARE, MANUAL, AND RELATED EQUIPMENT AND HARDWARE (WITH PROOF OF PAYMENT) TO THE PLACE OF PURCHASE FOR A FULL REFUND.

Cisco Systems, Inc. ("Cisco") and its suppliers grant to Customer ("Customer") a nonexclusive and nontransferable license to use the Cisco software ("Software") in object code form solely on a single central processing unit owned or leased by Customer or otherwise embedded in equipment provided by Cisco. Customer may make one (1) archival copy of the software provided Customer affixes to such copy all copyright, confidentiality, and proprietary notices that appear on the original. EXCEPT AS EXPRESSLY AUTHORIZED ABOVE, CUSTOMER SHALL NOT: COPY, IN WHOLE OR IN PART, SOFTWARE OR DOCUMENTATION; MODIFY THE SOFTWARE; REVERSE COMPILE OR REVERSE ASSEMBLE ALL OR ANY PORTION OF THE SOFTWARE; OR RENT, LEASE, DISTRIBUTE, SELL, OR CREATE DERIVATIVE WORKS OF THE SOFTWARE.

Customer agrees that aspects of the licensed materials, including the specific design and structure of individual programs, constitute trade secrets and/or copyrighted material of Cisco. Customer agrees not to disclose, provide, or otherwise make available such trade secrets or copyrighted material in any form to any third party without the prior written consent of Cisco. Customer agrees to implement reasonable security measures to protect such trade secrets and copyrighted material. Title to Software and documentation shall remain solely with Cisco.

This License is effective until terminated. Customer may terminate this License at any time by destroying all copies of Software including any documentation. This License will terminate immediately without notice from Cisco if Customer fails to comply with any provision of this License. Upon termination, Customer must destroy all copies of Software.

Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licenses to export, re-export, or import Software.

This License shall be governed by and construed in accordance with the laws of the State of California, United States of America, as if performed wholly within the state and without giving effect to the principles of conflict of law. If any portion hereof is found to be void or unenforceable, the remaining provisions of this License shall remain in full force and effect. This License constitutes the entire License between the parties with respect to the use of the Software.

Restricted Rights - Cisco's software is provided to non-DOD agencies with RESTRICTED RIGHTS and its supporting documentation is provided with LIMITED RIGHTS. Use, duplication, or disclosure by the U.S. Government is subject to the restrictions as set forth in subparagraph "C" of the Commercial Computer Software - Restricted Rights clause at FAR 52.227-19. In the event the sale is to a DOD agency, the U.S. Government's rights in software, supporting documentation, and technical data are governed by the restrictions in the Technical Data Commercial Items clause at DFARS 252.227-7015 and DFARS 227.7202.

LIMITED WARRANTY

Software. Cisco warrants that for a period of ninety (90) days from the date of shipment from Cisco: (i) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (ii) the Software substantially conforms to its published specifications. Except for the foregoing, the Software is provided AS IS. This limited warranty extends only to Customer as the original licensee. Customer's exclusive remedy and the entire liability of Cisco and its suppliers under this limited warranty will be, at Cisco or its service center's option, repair, replacement, or refund of the Software if reported (or, upon request, returned) to the party supplying the Software to Customer. In no event does Cisco warrant that the Software is error free or that Customer will be able to operate the Software without problems or interruptions.

Hardware. Cisco warrants that for a period of ninety (90) days from the date of shipment from Cisco that the Hardware will be free from defects in material and workmanship under normal use. This limited warranty extends only to Customer as original purchaser. Customer's exclusive remedy and the entire liability of Cisco and its suppliers under this limited warranty will be, at Cisco's or its service center's option, shipment of an advance replacement within five (5) working days at Cisco's expense, or a refund of the purchase price if the Hardware is returned to the party supplying it to Customer, freight and insurance prepaid. Cisco replacement parts used in Hardware repair may be new or equivalent to new.

Restrictions. This warranty does not apply if the product (a) has been altered, except by Cisco, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Cisco, (c) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident, or (d) is used in ultrahazardous activities.

DISCLAIMER OF WARRANTY. EXCEPT AS SPECIFIED IN THIS WARRANTY, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE, ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW.

IN NO EVENT WILL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR DATA, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY ARISING OUT OF THE USE OF OR INABILITY TO USE THE PRODUCT EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In no event shall Cisco's or its suppliers' liability to Customer, whether in contract, tort (including negligence), or otherwise, exceed the price paid by Customer. The foregoing limitations shall apply even if the above-stated warranty fails of its essential purpose.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

- Turn the television or radio antenna until the interference stops.
- Move the equipment to one side or the other of the television or radio.
- Move the equipment farther away from the television or radio.
- Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The following third-party software may be included with your product and will be subject to the software license agreement:

CiscoWorks software and documentation are based in part on HP OpenView under license from the Hewlett-Packard Company. HP OpenView is a trademark of the Hewlett-Packard Company. Copyright ©1992, 1993 Hewlett-Packard Company.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright ©1981, Regents of the University of California.

Network Time Protocol (NTP). Copyright ©1992, David L. Mills. The University of Delaware makes no representations about the suitability of this software for any purpose.

Point-to-Point Protocol. Copyright ©1989, Carnegie-Mellon University. All rights reserved. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

The Cisco implementation of TN3270 is an adaptation of the TN3270, curses, and termcap programs developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981-1988, Regents of the University of California.

Cisco incorporates Fastmac and TrueView software and the RingRunner chip in some Token Ring products. Fastmac software is licensed to Cisco by Madge Networks Limited, and the RingRunner chip is licensed to Cisco by Madge NV. Fastmac, RingRunner, and TrueView are trademarks and in some jurisdictions registered trademarks of Madge Networks Limited. Copyright ©1995, Madge Networks Limited. All rights reserved.

XRemote is a trademark of Network Computing Devices, Inc. Copyright © 1989, Network Computing Devices, Inc., Mountain View, California. NCD makes no representations about the suitability of this software for any purpose.

The X Window System is a trademark of the X Consortium, Cambridge, Massachusetts. All rights reserved.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AtmDirector, AutoConnect, AutoRoute, AXIS, BPX, Catalyst, CD-PAC, CiscoAdvantage, CiscoFusion, Cisco IOS, the Cisco IOS logo, *CiscoLink*, CiscoPro, the CiscoPro logo, CiscoRemote, the CiscoRemote logo, CiscoSecure, Cisco Systems, CiscoView, CiscoVision, CiscoWorks, ClickStart, ControlStream, EdgeConnect, EtherChannel, FairShare, FastCell, FastForward, FastManager, FastMate, FastPADImp, FastPADmicro, FastPADmp, FragmentFree, FrameClass, Fulcrum INS, IGX, Impact, Internet Junction, JumpStart, LAN^{2}LAN Enterprise, LAN^{2}LAN Remote Office, LightSwitch, NetBeyond, NetFlow, Newport Systems Solutions, *Packet*, PIX, Point and Click Internetworking, RouteStream, Secure/IP, SMARTnet, StrataSphere, StrataSphere BILLder, StrataSphere Connection Manager, StrataSphere Modeler, StrataSphere Optimizer, Stratm, StrataView Plus, StreamView, SwitchProbe, SwitchVision, SwitchWare, SynchroniCD, *The Cell*, The FastPacket Company, TokenSwitch, TrafficDirector, Virtual EtherSwitch, VirtualStream, VlanDirector, Web Clusters, WNIC, Workgroup Director, Workgroup Stack, and XCI are trademarks; Access by Cisco, Bringing the Power of Internetworking to Everyone, Enter the Net with MultiNet, and The Network Works. No Excuses. are service marks; and Cisco, the Cisco Systems logo, CollisionFree, Combinet, EtherSwitch, FastHub, FastLink, FastNIC, FastPacket, FastPAD, FastSwitch, ForeSight, Grand, Grand Junction, Grand Junction Networks, the Grand Junction Networks logo, HSSI, IGRP, IPX, Kalpana, the Kalpana logo, LightStream, MultiNet, MultiWare, OptiClass, Personal Ethernet, Phase/IP, RPS, StrataCom, TGV, the TGV logo, and UniverCD are registered trademarks of Cisco Systems, Inc. All other trademarks, service marks, registered trademarks, or registered service marks mentioned in this document are the property of their respective owners.

Introduction to Cisco Router Configuration: Student Guide
Copyright © 1996, Cisco Systems, Inc.
All rights reserved. Printed in USA.
969R

TABLE OF CONTENTS

Module 2 Getting Started with Cisco IOS Software

Module 3 Networking Protocol Suites

Copyright © 1996, Cisco Systems, Inc.

Copyright © 1996, Cisco Systems, Inc.

Course Introduction

Copyright © 1996, Cisco Systems, Inc.

Curriculum Related to ICRC/ACRC

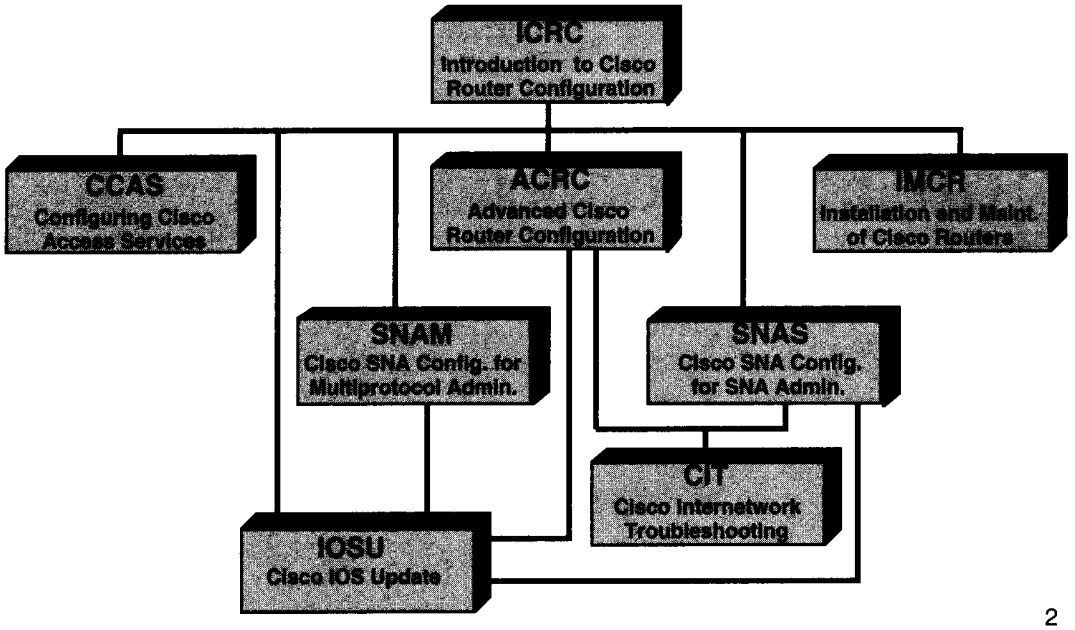

2

Copyright © 1996, Cisco Systems, Inc.

Course Prerequisites

- **Basic knowledge of:**
  - **— Networking**
  - **— Binary and hexadecimal numbering (recommended)**

3

Copyright © 1996, Cisco Systems, Inc.

Welcome

- Name and company
- Job title and responsibilities
- Internetworking experience
- Prerequisites
- Your objectives

4

Copyright © 1996, Cisco Systems, Inc.

Housekeeping

- Sign-in sheet
- Class and break times
- Attire
- Student materials

5

Copyright © 1996, Cisco Systems, Inc.

Housekeeping (cont.)

- Questions in class
- Off-line questions

  This class covers announced products only

  We cannot discuss information covered by nondisclosure restrictions

6

Copyright © 1996, Cisco Systems, Inc.

Housekeeping (cont.)

- Break/lunch facilities
- Restrooms
- Telephones
- E-mail connections
- Messages
- Emergency procedures
- Smoking policy

7

Copyright © 1996, Cisco Systems, Inc.

Course Objective

- **Upon completion of this course, you will be able to configure Cisco routers for operation in multiprotocol internetworks**

8

Copyright © 1996, Cisco Systems, Inc.

Course Agenda

- Day 1
  - Class Startup routines
  - Introduction to Internetworking
  - Getting Started with Cisco IOS
- Day 2
  - Getting Started with Cisco IOS (cont.)
  - Networking Protocol Suites (TCP/IP)

9

Copyright © 1996, Cisco Systems, Inc.

Course Agenda (cont.)

- Day 3
  - Networking Protocol Suites (TCP/IP cont.)
  - Networking Protocol Suites (Novell IPX)
- Day 4
  - Networking Protocol Suites (AppleTalk)
  - Networking Protocol Suites (Access Lists)
  - Wide-Area Networking (PPP, DDR, ISDN)
- Day 5
  - Wide-Area Networking (X.25, Frame Relay, AutoInstall)

10

Copyright © 1996, Cisco Systems, Inc.

Graphic Symbols

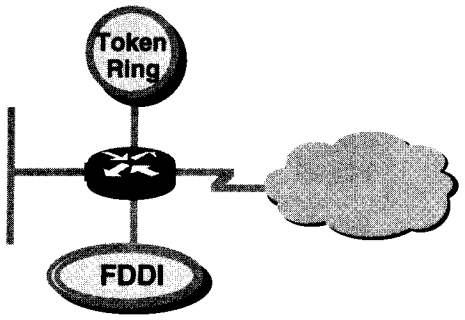

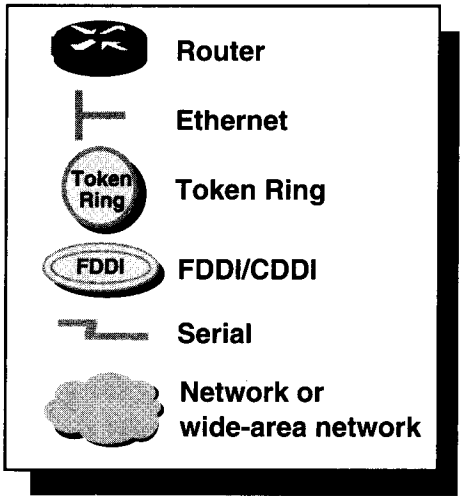

11

Copyright © 1996, Cisco Systems, Inc.

Graphic Symbols (cont.)

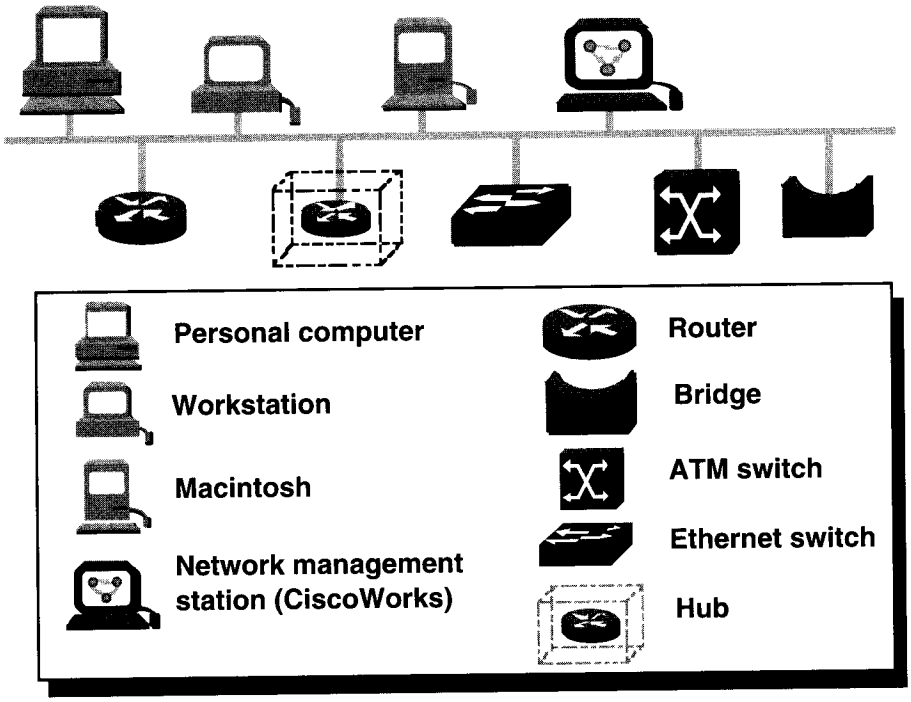

12

Copyright © 1996, Cisco Systems, Inc.

MODULE 1

Introduction to Internetworking

CHAPTER 1

The Internetworking Model

Copyright © 1996, Cisco Systems, Inc.

Objectives

Upon completion of this chapter, you will be able to:

Discuss the major influences of user requirements on network evolution

Identify at least three reasons why the industry uses a layered network model

Identify the functions of each layer of the ISO/OSI reference model

Define and explain the five conversion steps of data encapsulation

2

This chapter discusses how networks have evolved and the improvements made to the services that are available to network users. It also discusses the layered International Organization for Standardization/Open Systems Interconnect (ISO/OSI) reference model and explains how data is encapsulated for transmission according to the model.

Sections:

- Network Evolution
- The Layered Model
- Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

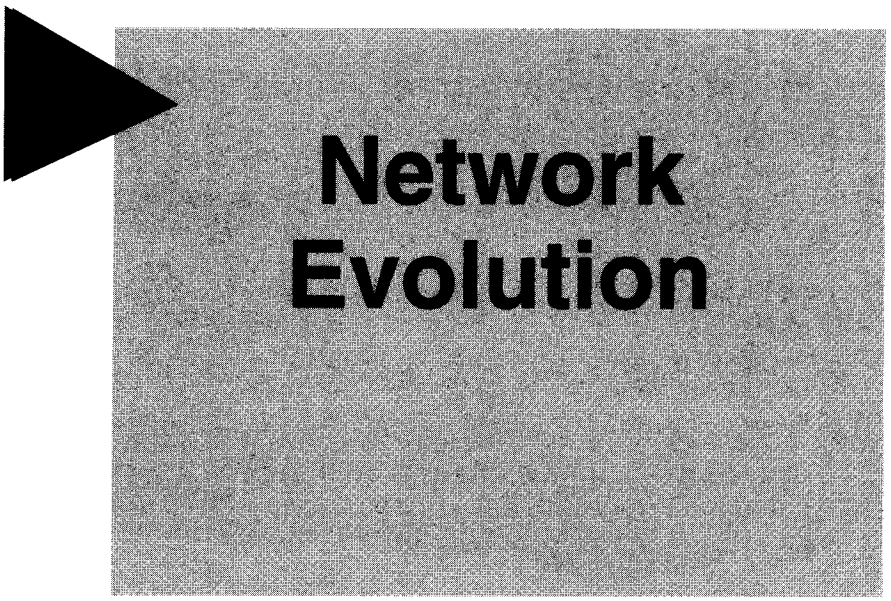

3

Network Evolution

Copyright © 1996, Cisco Systems, Inc.

1960s and 1970s: Communications

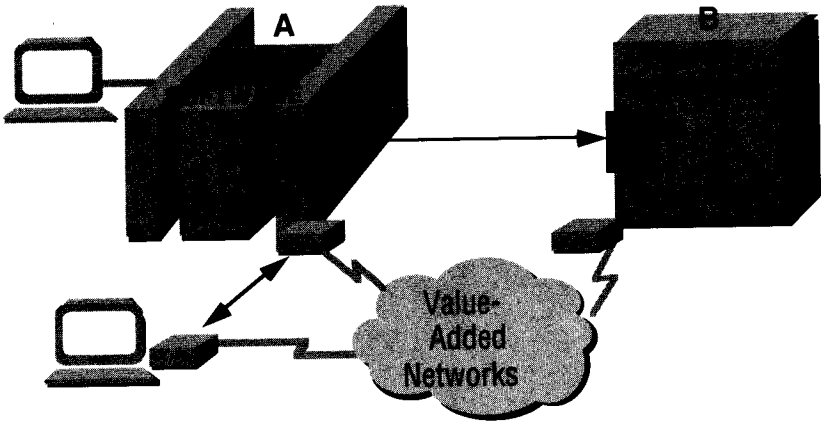

4

In the 1960s and early 1970s, the traditional computer communication environment centered around the host (mainframe). The centralized computing environment of the past required low-speed access lines that unintelligent terminals used to communicate to the centralized host. IBM computers with Systems Network Architecture (SNA) networks and non-IBM computers with X.25 public data networks are typical examples of this type of environment. The graphic illustrates a simple host-based communication environment.

On a single computer, accessing resources, running programs, and copying files are relatively straightforward. The computer must identify the requesting user and the desired destination device or program and then coordinate access between them. The single computer in this scenario is the master of all resources and thus can easily manage and coordinate them.

Even in a network of only two computers, coordinating resources becomes much more complex. Transferring information requires, among other things, addressing, error detection, error correction, synchronization, and transmission coordination.

Copyright © 1996, Cisco Systems, Inc.

1970s and 1980s: Networks

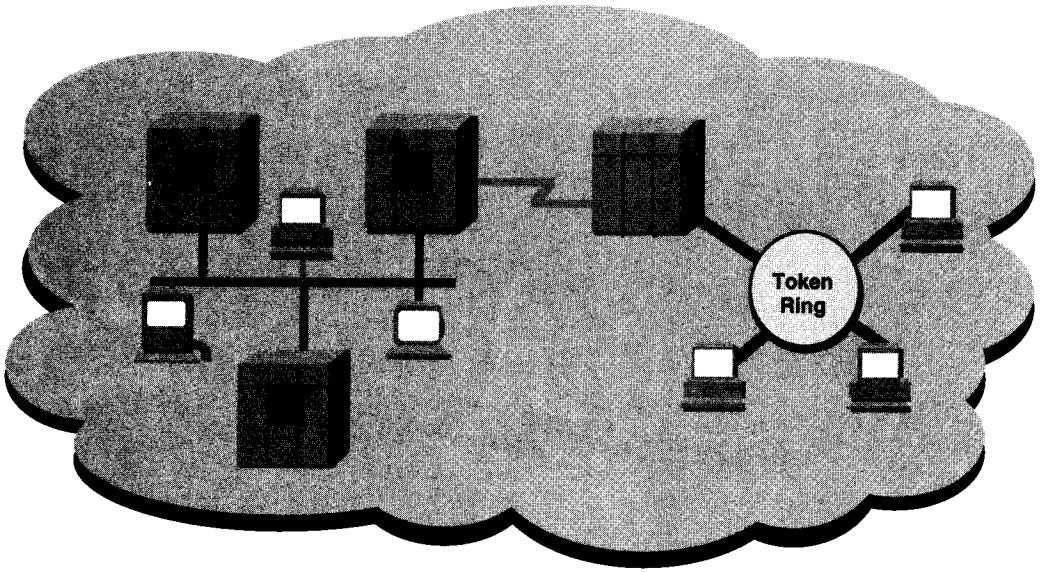

5

The introduction of PCs revolutionized traditional communication and computer networks. Initially, PCs were standalone devices. As businesses realized the flexibility and the power of these devices, their use increased. Local-area networks (LANs) evolved primarily to decrease the cost of expensive devices such as printers and hard disks.

The strategic importance of interconnected networks was quickly realized. Organizations began to move toward linking previously isolated LANs. Interconnected networks provided the basis for enterprise-wide applications such as e-mail and file transfer. These in turn increased overall productivity and competitiveness.

In the 1970s and 1980s, minicomputers and shared wide-area networks (WANs) evolved. Minicomputers, often located away from the central data center, facilitated the emergence of distributed data processing. The Digital Equipment Corporation VAX systems and DECnet networking are typical of this era. In general, however, applications remained separate and independent, and different communication protocols were developed.

Copyright © 1996, Cisco Systems, Inc.

1980s and 1990s: Internetworks

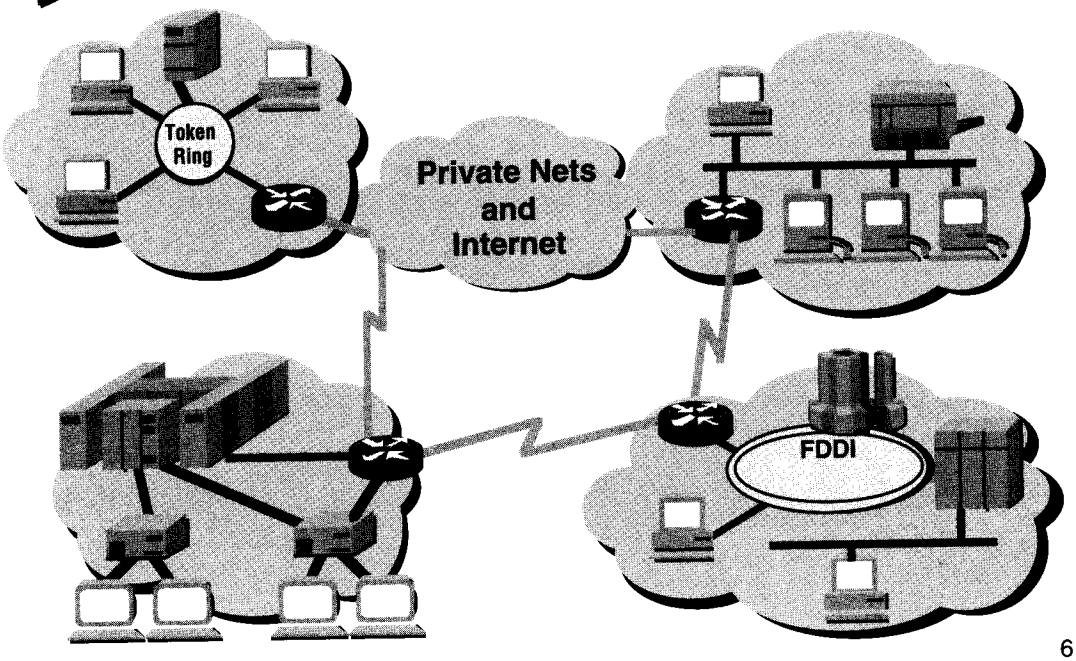

6

For many businesses, today's networks are a mixture of old and new technologies. IBM networks might be operating virtually in parallel with the newer LAN interconnected networks, electronic commerce, and messaging systems. Local networks, public data networks, leased lines, and high-speed mainframe channels have all been used on an opportunistic basis with little regard for overall integration and consistency. Application migrating from central hosts to distributed servers has resulted in new networking requirements and changing traffic patterns.

The approach to computer communication in most organizations is changing rapidly in response to new technologies, evolving business requirements, and the need for "instant" knowledge transfer. To meet these requirements, the internetwork, whatever form it takes, must be flexible, scalable, and adaptable to suit any organizational level (branch, regional, headquarters).

Internetworks tie LANs and WANs, computer systems, software, and related devices together to form the corporate communication infrastructure. An internetwork moves information anywhere within a corporation and to external suppliers and customers. By serving as the organization's information highway, the internetwork has become a key strategic asset and a competitive advantage.

Copyright © 1996, Cisco Systems, Inc.

1990s: Global Internetworking

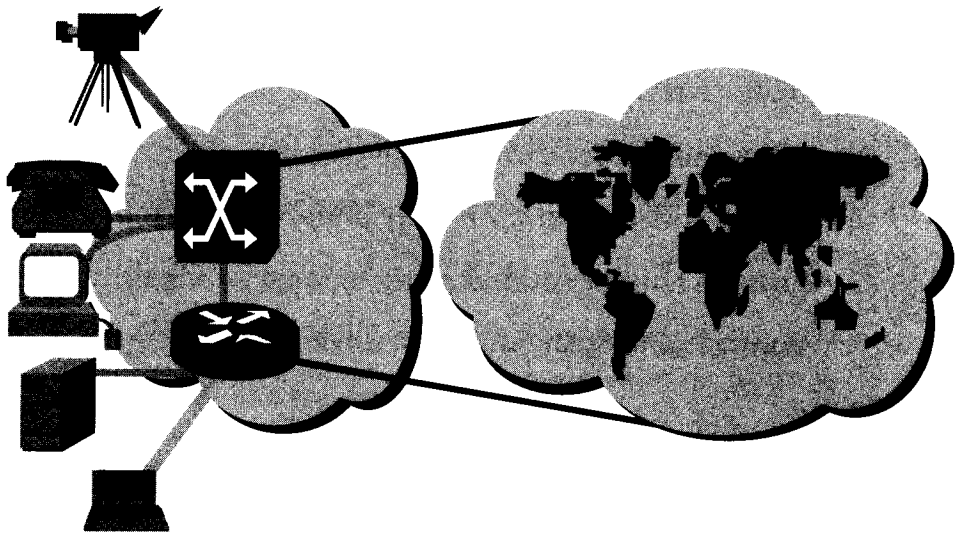

7

Studies show that users increasingly require more bandwidth. Networks will have to meet these demands as well as provide low delay, bandwidth on demand, and other new services. New devices will take their place alongside the router as additional network tools. Current and future networks will have more functions distributed and must provide for the integration of voice, data, and video. Such networks are characterized by the following:

- Increasing use of graphics and imaging
- Larger files
- Larger programs
- Client/server computing
- Bursty network traffic

Global internetworking will provide an environment for emerging applications that will require even greater amounts of bandwidth. Many of these applications are driven by the evolution of multimedia requirements that have a high-definition image, full-motion video, or a digitized audio component.

Copyright © 1996, Cisco Systems, Inc.

Local-Area Networks and Devices

LANs are designed to:

- **Operate within a limited geographic area**
- **Allow multiaccess to high-bandwidth media**
- **Control the network privately under local administration**
- **Provide full-time connectivity to local services**
- **Connect physically adjacent devices**

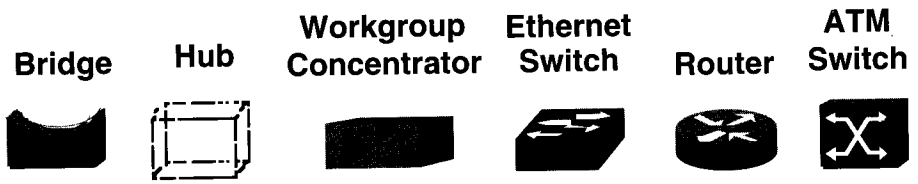

8

Major characteristics of LANs are:

- The network operates within a building or floor of a building. The geographic scope for ever-more powerful LAN desktop devices running more powerful applications is for less area per LAN.
- LANs provide multiple connected desktop devices (usually PCs) with access to high-bandwidth media.
- An enterprise purchases the media and connections used in the LAN; the enterprise can privately control the LAN as it chooses.
- Local services are usually available; LANs rarely shut down or restrict access to connected workstations.
- By definition, the LAN connects physically adjacent devices on the media. LAN devices include:
  - Bridges that connect LAN segments and help filter traffic.
  - Hubs that concentrate LAN connection and allow use of twisted-pair copper media.
  - Workgroup concentrators that deliver 100-Mbps service over fiber or copper cabling.
  - Ethernet switches that offer full-duplex, dedicated bandwidth to segments or desktops.
  - Routers that offer many services including internetworking and broadcast control.
  - Asynchronous Transfer Mode (ATM) switches that provide high-speed cell switching.

Copyright © 1996, Cisco Systems, Inc.

Wide-Area Networks and Devices

WANs are designed to:

- **Operate over geography of telecommunication carriers**
- **Allow access over serial interfaces operating at lower speeds**
- **Control the network subject to regulated public services**
- **Provide full-time and part-time connectivity**
- **Connect devices separated over wide, even global areas**

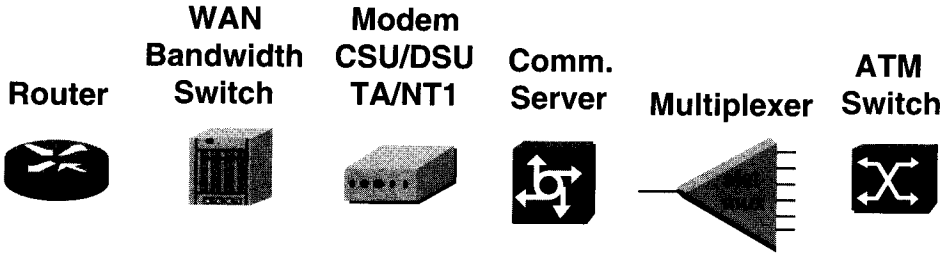

9

Major characteristics of WANs are:

- The network operates beyond the local LAN's geographic scope. It uses the services of carriers such as Regional Bell Operating Companies (RBOCs), Sprint, and MCI.
- WANs use serial connections of various types to access bandwidth over wide-area geographies.
- An enterprise pays the carrier or service provider for connections used in the WAN; the enterprise can choose which services it uses. Carriers are usually regulated by tariffs.
- WANs rarely shut down, but because the enterprise must pay for services used, it might restrict access to connected workstations. All WAN services are not available in all locations.
- By definition, the WAN connects devices separated by wide areas. WAN devices include:
  - — Routers that offer many services including internetworking and WAN interface controls.
  - — Switches that connect to WAN bandwidth for X.25, Frame Relay, and voice, data, and video communication. These WAN switches can share bandwidth among allocated service priorities, recover from outages, and provide network design and management systems.
  - — Modems that interface voice-grade services; channel service units/digital service units (CSU/DSU) that interface T1/E1 services; Terminal Adapters/Network Termination 1 (TA/NT1) that interface Integrated Services Digital Network (ISDN) services.
  - — Communication servers that concentrate dialin and dial-out user communication.
  - — Multiplexers that share a WAN facility among several demand channels.
  - — ATM switches that provide high-speed cell switching.

Copyright © 1996, Cisco Systems, Inc.

Enterprise Developments

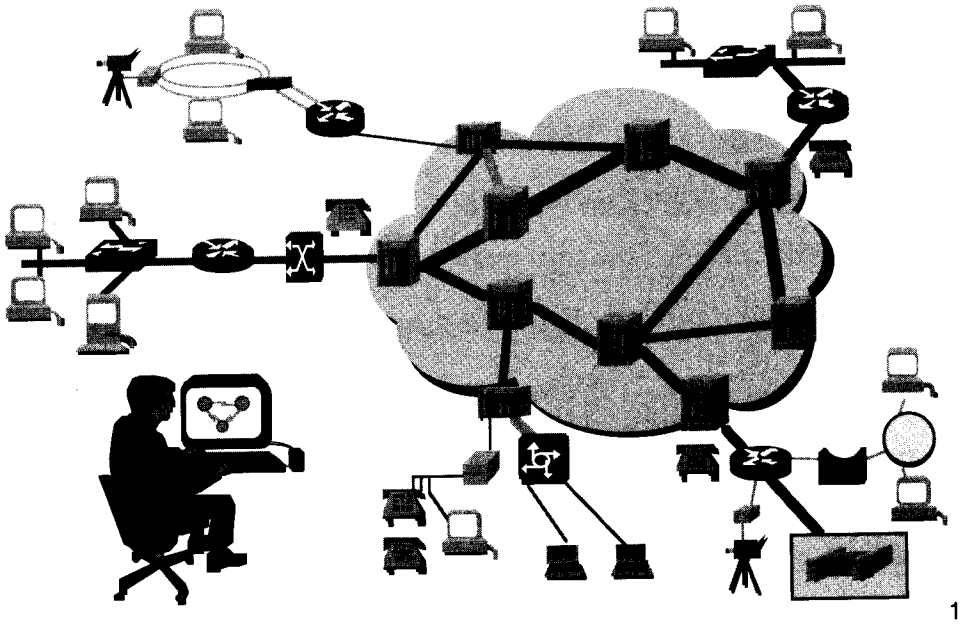

10

The enterprise is a corporation, agency, service provider, or other organization that will tie together its data, communication, computing, and storage resources. An enterprise network usually contains a hybrid of both private and public network elements. Any or all of the LAN and WAN devices shown on the previous pages can be found on the enterprise network.

Developments on the enterprise network include:

- LANs interconnected to provide client/server applications integrated with the traditional legacy applications from mainframe data centers.
- End-user needs for higher bandwidth on the LANs, which can be consolidated at a switch and delivered on dedicated media.
- Integration of formerly separate networks so that the nonbursty traffic from voice and video applications coexist on a single network.
- Relaying technologies for WAN service, with very rapid growth in Frame Relay and more gradual growth of cell relay (for example, ATM)

Cisco was the first company to offer products that worked from the desktop all the way to the central office switch of the telecommunication carriers. With the StrataCom acquisition, Cisco adds the missing pieces that work inside the WAN cloud. Now Cisco is able to offer products that cover the entire realm for enterprise network development.

Copyright © 1996, Cisco Systems, Inc.

Network Manager Tasks and Goals

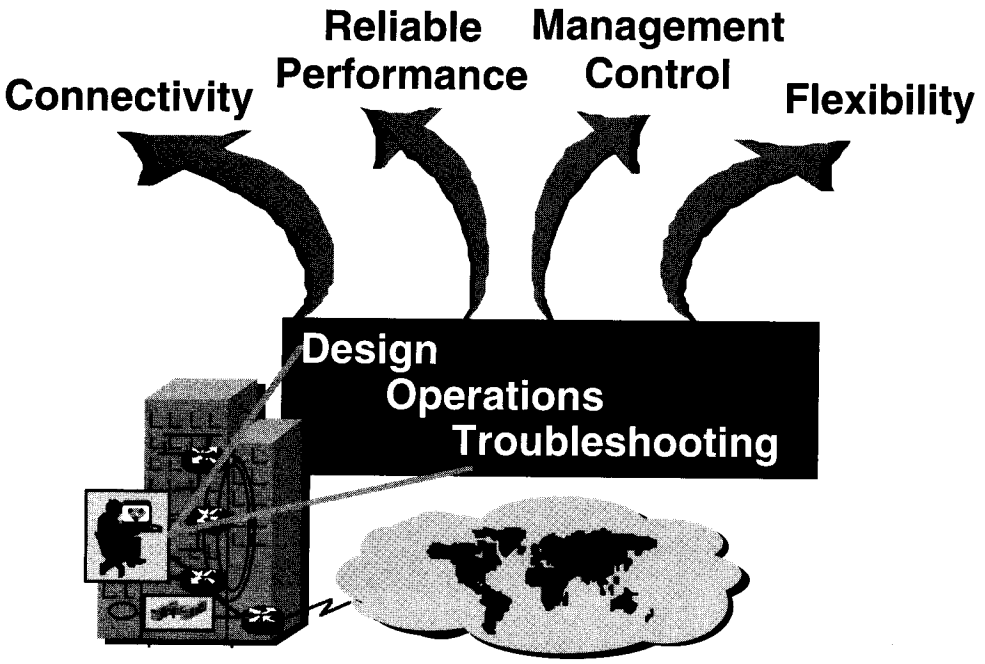

11

The tasks and goals of modern network managers center around four major categories:

- Connectivity—The internetwork must serve those in the organization who depend on it. Regardless of the range of media attachments, transmission speeds, and other technical details, the network design connects previously separate resources.
- Reliable performance—The organization becomes increasingly dependent on its internetworking tools including the operator interface, the ability to distribute network software updates, utilities to log and monitor performance, and the functions to secure access to resources.
- Management control—An internetwork provides crucial functions; it also expends critical resources. Administrators continually ask how they can improve management controls and security. After the network is designed and operational, troubleshooting tasks follow.
- Flexibility—Expanding internetworks demand administrator flexibility. Expansion and consolidation efforts mean overcoming physical or geographic boundaries. Enterprises seek ways to provide new services and products to a network-accessible, global economy.

Copyright © 1996, Cisco Systems, Inc.

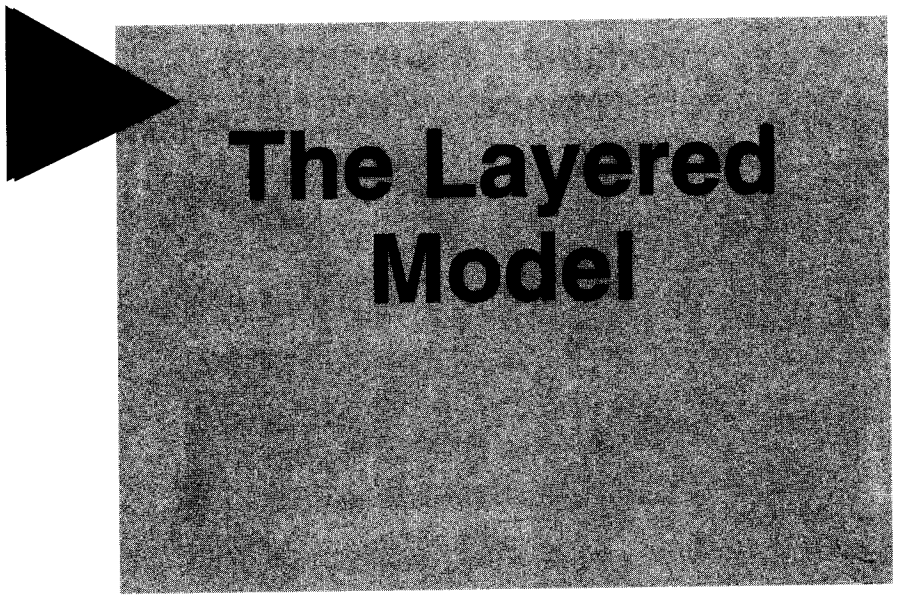

12

The Layered Model

Copyright © 1996, Cisco Systems, Inc.

Why a Layered Network Model?

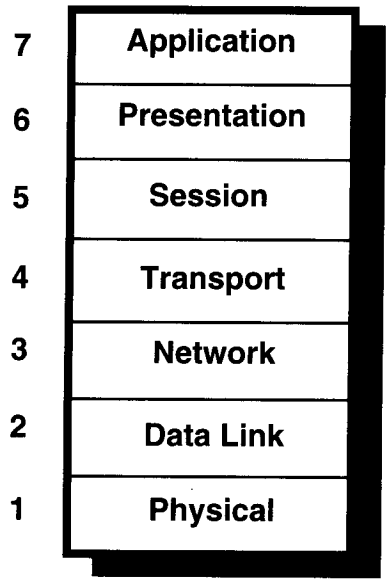

- **Reduces complexity**
- **Standardizes interfaces**
- **Facilitates modular engineering**
- **Ensures interoperable technology**
- **Accelerates evolution**
- **Simplifies teaching and learning**

13

Most communication environments separate the communication functions and application processing. This separation of networking functions is called layering. For the OSI model, seven numbered layers indicate distinct functions. Within the Transmission Control Protocol/Internet Protocol (TCP/IP), for example, distinct functions fit into five named layers. Regardless of the number of layers, the reasons for this division of network functions include the following:

- Divide the interrelated aspects of network operation into less complex elements.
- Define standard interfaces for plug-and-play compatibility and multivendor integration.
- Enable engineers to specialize design and development efforts on modular functions.
- Promote symmetry in the different internetwork modular functions so they interoperate.
- Prevent changes in one area from impacting other areas, so each area can evolve more quickly.
- Divide the complexity of internetworking into discrete, more easily learned operation subsets.

Note A layered model does not define or constrain an implementation; it provides a framework. Implementations, therefore, do not conform to the OSI reference model, but they do conform to the standards developed from the OSI reference model principles.

Copyright © 1996, Cisco Systems, Inc.

Layer Functions

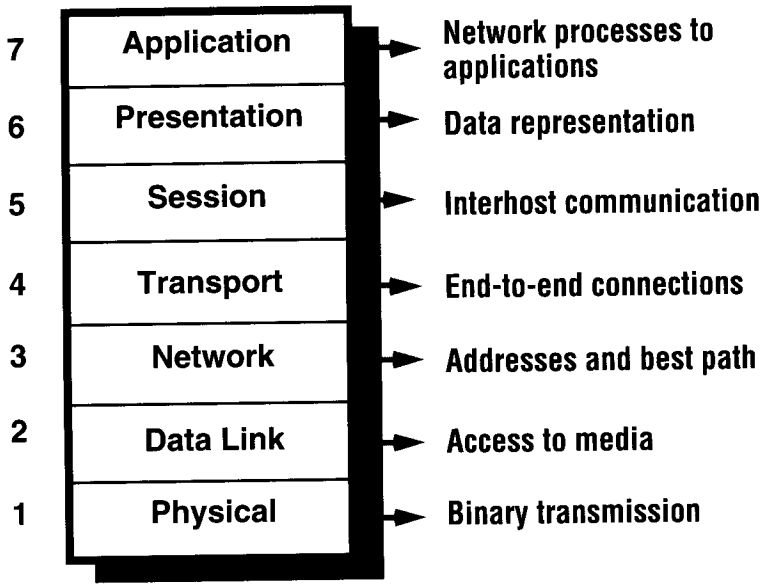

14

Each layer of the ISO model serves a specific function. Those functions are defined by the OSI) and can be used by any network products vendor. The functions are:

- Application—The application layer provides network services to user applications. For example, a word processing application is serviced by file transfer services at this layer.
- Presentation—This layer provides data representation and code formatting. It ensures that the data that arrives from the network can be used by the application, and it ensures that information sent by the application can be transmitted on the network.
- Session—This layer establishes, maintains, and manages sessions between applications.
- Transport—This layer segments and reassembles data into a data stream.
- Network—This layer determines the best way to move data from one place to another. It manages device addressing and tracks the location of devices on the network. The router operates at this layer.
- Data Link—This layer provides physical transmission across the medium. It handles error notification, network topology, and flow control.
- Physical—This layer provides the electrical, mechanical, procedural, and functional means for activating and maintaining the physical link between systems.

Copyright © 1996, Cisco Systems, Inc.

Peer-to-Peer Communication

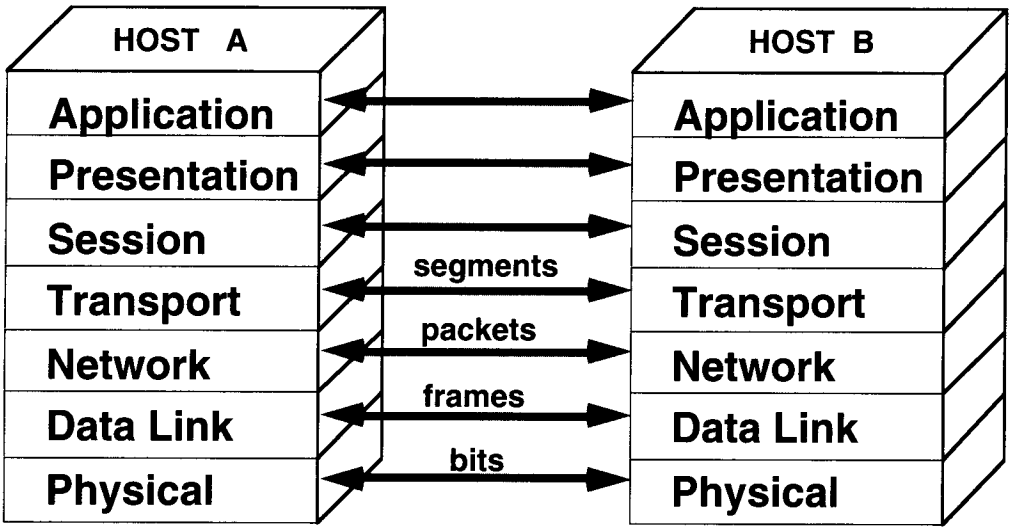

15

Each layer uses its own layer protocol to communicate with its peer layer in the other system. Each layer's protocol exchanges information, called protocol data units (PDUs), between peer layers. A given layer can use a more specific name for its PDU.

For example, in TCP/IP the transport layer of TCP communicates with the peer TCP function using segments.

This peer-layer protocol communication is achieved by using the services of the layers below it. The layer below any current layer provides its services to the current layer. Each lower-layer service takes upper-layer information as part of the lower-layer PDUs it exchanges with its layer peer.

Thus, the TCP segments become part of the network layer packets (also called datagrams) exchanged between IP peers. In turn, the IP packets must become part of the data link frames exchanged between directly connected devices. Ultimately, these frames must become bits as the data is finally transmitted by the physical-layer protocol using hardware.

Copyright © 1996, Cisco Systems, Inc.

Data Encapsulation

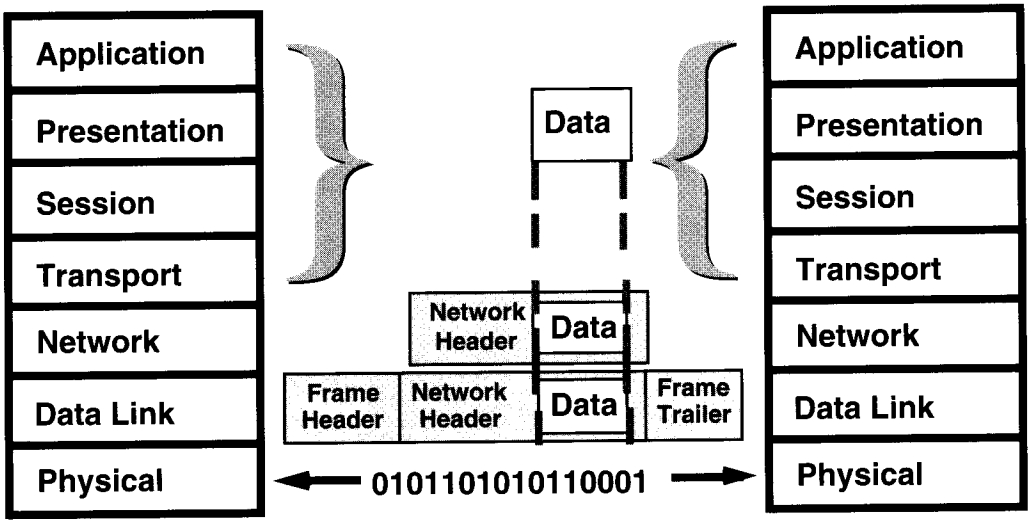

16

Each layer depends on the service function of the ISO/OSI layer below it. To provide this service, the lower layer uses encapsulation to put the PDU from the upper layer into its data field; then it can add whatever headers and trailers the layer will use to perform its function.

For example, the network layer provides a service to the transport layer, and the transport layer presents "data" to the internetwork subsystem.

The network layer has the task of moving that data through the internetwork. It accomplishes this task by encapsulating the data within a header. This header contains information required to complete the transfer, such as source and destination logical addresses.

The data link layer in turn provides a service to the network layer. It encapsulates the network layer information in a frame. The frame header contains information required to complete the data link functions. For example, the frame header contains physical addresses.

The physical layer also provides a service to the data link layer. This service includes encoding the data link frame into a pattern of ones and zeros for transmission on the medium (usually a wire).

Copyright © 1996, Cisco Systems, Inc.

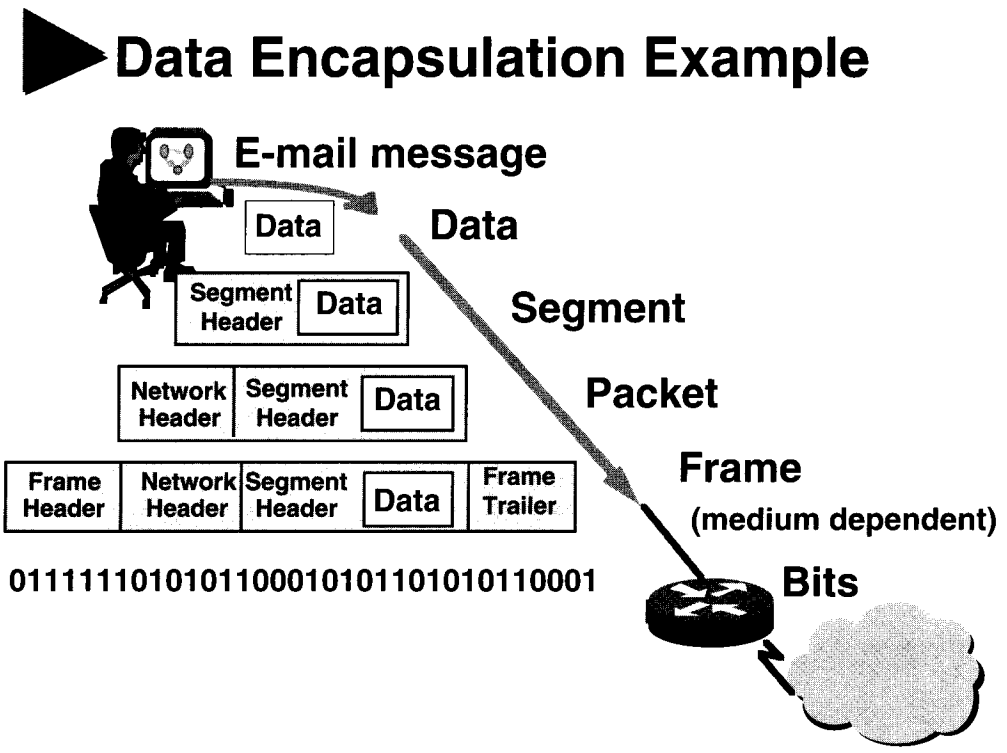

As internetworks perform services for users, the flow and packaging of the information changes. In this example of internetworking, five conversion steps occur:

Step 1 As a user sends an e-mail message, its alphanumeric characters are converted to use the internetwork. This is the data.

Step 2 One change packages the message "data" for the internetwork transport subsystem. By using segments, the transport function ensures that the message hosts at both ends of the e-mail system can reliably communicate.

Step 3 The next change prepares the data so they can use the internetwork by putting the data into a packet or datagram that contains a network header with source and destination logical addresses. These addresses help network devices send the packets across the network along a chosen path.

Step 4 Each network device must put the packet into a frame so it can communicate over its interface to the network. The frame allows connection to the next directly connected network device on the link. Each device in the chosen network path requires framing to connect to the next device.

Step 5 The frame must be converted into a pattern of 1s and 0s for transmission on the medium (usually a wire). Some clocking function enables the devices to distinguish these bits as they traverse the medium.

The medium on the physical internetwork can vary along the path used. For example, the e-mail message can originate on a LAN, cross a campus backbone, go out a low-speed WAN link, and use a higher-speed WAN link until it reaches its destination on another remote LAN.

Copyright © 1996, Cisco Systems, Inc.

Remaining Chapter Sequence

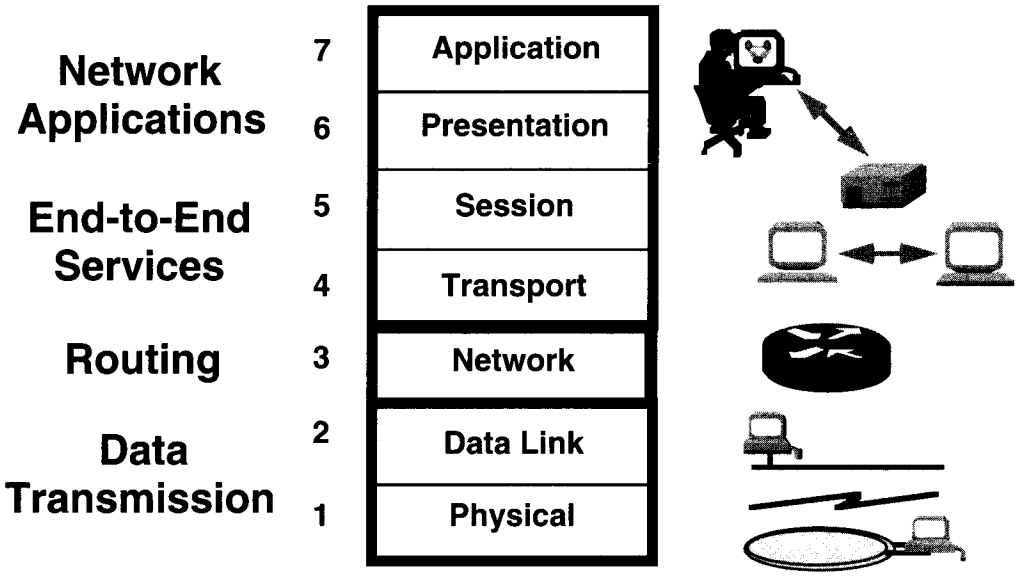

18

Now you have reviewed the evolution leading to the modern networks. You have seen the use of a model and have been introduced to the operations and functions at each layer. The remaining three chapters of this "Introduction to Internetworking" (I2I) module will proceed as follows:

- "Applications and Upper Layers"—Network applications layers and how they provide application, data presentation, and session functions; also the upper layer that provides end-to-end services between hosts using transport layer services.
- "Physical and Data Link Layers"—Data transmission services provided by lower-layer functions with specific variations for LAN and WAN framing and media.
- "Network Layer and Path Determination"—Routing using Layer 3 services of the network layer and other processes; Layer 3 is the primary domain of the router.

Copyright © 1996, Cisco Systems, Inc.

Summary

Internetworking evolves to support current and future applications

The OSI reference model organizes network functions into seven categories called layers

Data flows from upper-level user applications to lower-level bits transmitted over network media

Peer-to-peer functions use encapsulation and de-encapsulation at layer interfaces

Most network manager tasks configure the lower three layers

Copyright © 1996, Cisco Systems, Inc.

Exercise: The Internetworking Model

Problem 1

Objective: Identify at least three reasons why the industry uses a layered network model.

Write at least three reasons why the industry uses a layered network model.

Copyright © 1996, Cisco Systems, Inc.

Problem 2

Objective: Identify the functions of each layer of the ISO/OSI reference model.

As your instructor reads a layer's function in the ISO/OSI reference model, write that layer's name in the appropriate numbered box in the seven-layer rectangle below.

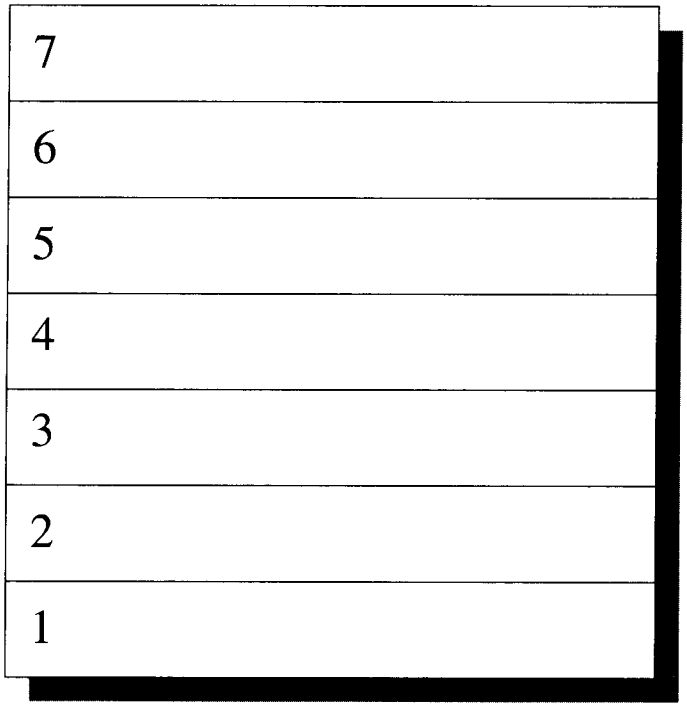

Copyright © 1996, Cisco Systems, Inc.

Problem 3

Objective: Define and explain the five conversion steps of data encapsulation.

When a user's information is converted into bits that hardware can use, five conversion steps can occur. Write these five steps on the lines provided below; include a brief description of what occurs during each conversion step. The first and last steps are provided as your reference.

A) User information is converted to_____.

B) _____is converted to _____

C) _____are converted to _____

D) _____are converted to _____

E) _____are converted to bits.

Copyright © 1996, Cisco Systems, Inc.

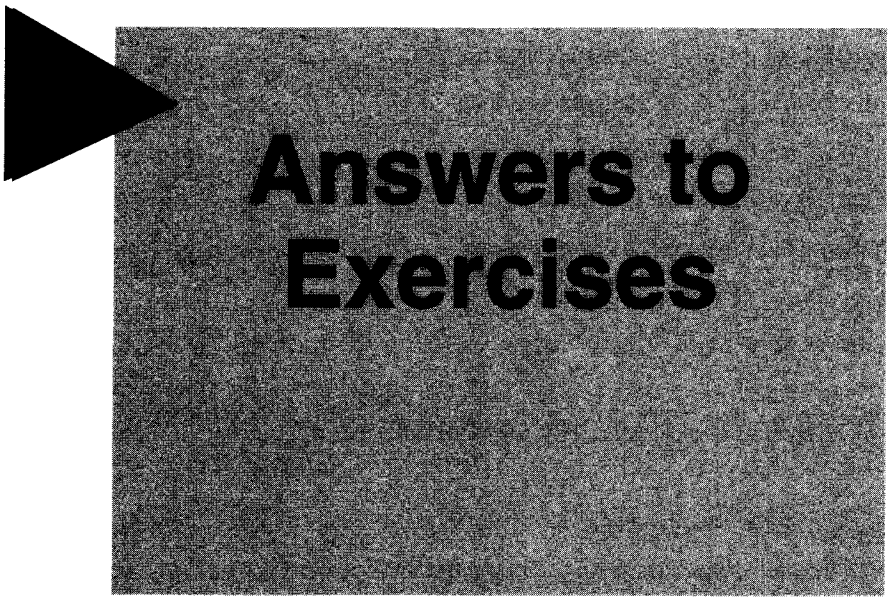

23

Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

Exercise: The Internetworking Model

Problem 1

- Clarify what general function is to be done rather than how to do it.
- Reduce the complexity of networking into more manageable sublayers.
- Enable interoperability (in theory at least) using standard interfaces.
- Allow changes in one layer to occur without changing other layers.
- Speed up network industry progress by allowing specialization.
- Allow for shortcut explanations to facilitate protocol comparisons.
- Order network troubleshooting steps (for example, check physical layer first).
- Facilitate systematic troubleshooting.

See the student page titled "Why a Layered Network Model?" for more information.

Problem 2

7 Application—Provides network or internetwork processes to computer applications.

6 Presentation—Defines data representation and code formatting.

5 Session—Synchronizes communication between applications on different hosts.

4 Transport—Provides end-to-end connections; can offer reliable transmission.

3 Network—Defines network addressing and determines the best path through an internetwork.

2 Data Link—Controls access to communication media between directly connected devices.

1 Physical—Sends and receives binary information using device interfaces.

Problem 3

A) User information is converted to data.

B) Data is converted to segments.

C) Segments are converted to packets or datagrams.

D) Packets or datagrams are converted to frames.

E) Frames are converted to bits.

Copyright © 1996, Cisco Systems, Inc.

CHAPTER 2

Applications and Upper Layers

Copyright © 1996, Cisco Systems, Inc.

Objectives

Upon completion of this chapter, you will be able to:

Name and describe computer, network, and internetwork applications

Describe the OSI presentation functions and identify common standards

Describe the OSI session functions and identify common standards

Describe the OSI transport functions for end-to-end network services

Identify common processes for establishing connections, flow control, and windowing

2

This chapter discusses the upper four layers of the OSI reference model: application, presentation, session, and transport. It briefly explains the function of the application, presentation, and session layers. The transport layer is covered in more detail, explaining how data is transmitted between the sender and the receiver.

Sections:

- Application, Presentation, and Session Layers
- Transport Layer
- Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

3

Application, Presentation, and Session Layers

Copyright © 1996, Cisco Systems, Inc.

Application Layer

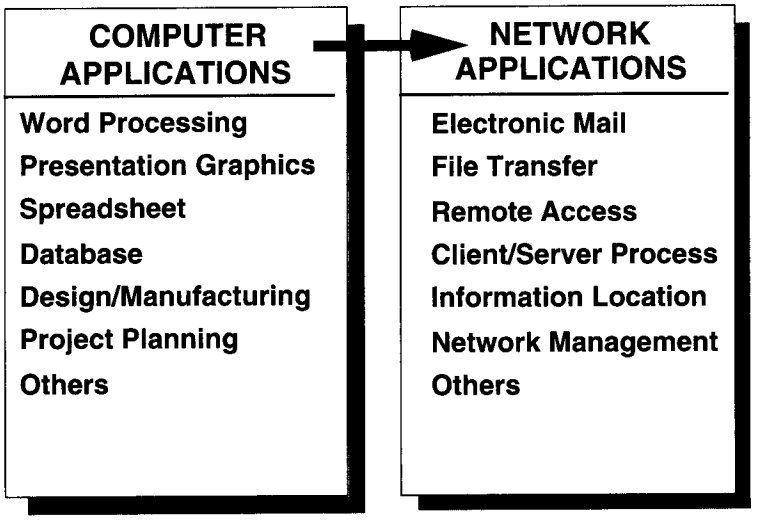

- **Selects network application to support user's application**

4

In the context of the OSI reference model, the application layer supports the communicating component of an application.

Computer applications can require only desktop resources. However, an application might incorporate a communicating component from one or more network applications. Several types of network applications are listed in the right column, Network Applications.

An application must have a communicating component to be relevant to a discussion of internetworking. For example, a word processor might incorporate a file transfer component that allows a document to be transferred electronically over telecommunication facilities. This file transfer component qualifies the word processor as an application in the OSI context and belongs in Layer 7 of the OSI reference model.

Copyright © 1996, Cisco Systems, Inc.

Application Layer (cont.)

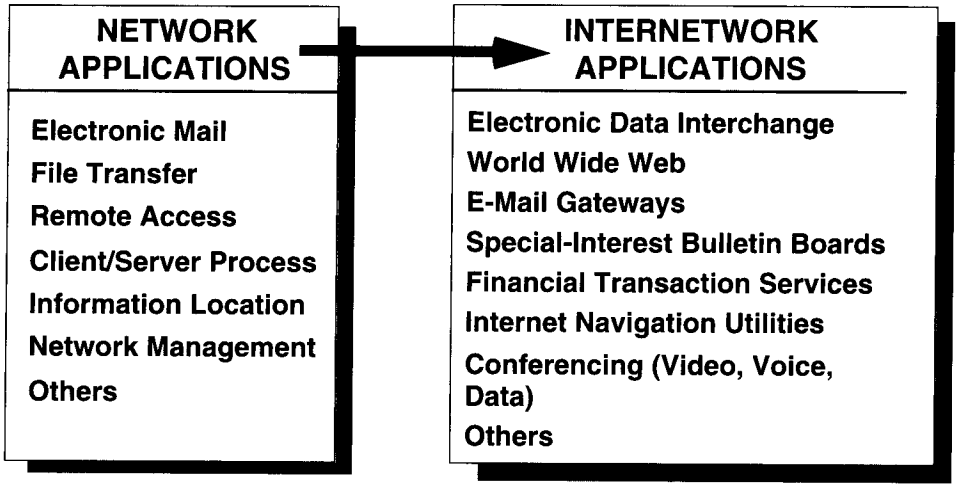

- **Internetwork applications can extend beyond the enterprise**

5

Many of the network applications offer services for enterprise communication. However, a growing requirement for internetworking in the 1990s and later extends beyond the enterprise. Information exchanges and commerce between organizations increasingly involve internetworking applications such as those listed in the right column of the graphic.

- Electronic data interchange (EDI) offers specialized standards and processes to improve the flow of orders, shipments, inventories, and accounting between businesses.
- The World Wide Web (WWW) links thousands of servers using a variety of formats including text, graphics, video, and sound. Browsers such as Mosaic and Netscape simplify access and viewing.
- The e-mail gateways might use the X.400 standard or Simple Mail Transfer Protocol (SMTP) to pass messages between different e-mail applications.
- Thousands of special-interest bulletin boards connect people who can chat with each other, post messages, and share public-domain software.
- Transaction services aimed at the financial community obtain and sell information including investment, market, commodity, currency, and credit data to subscribers.
- Special-purpose applications such as Gopher, Fetch, and Wide Area Information Server (WAIS) help navigate the way to resources on the Internet.
- People located in different regions use conferencing applications to communicate with live and prefilmed video, voice, data, and fax exchange.

Copyright © 1996, Cisco Systems, Inc.

Presentation Layer

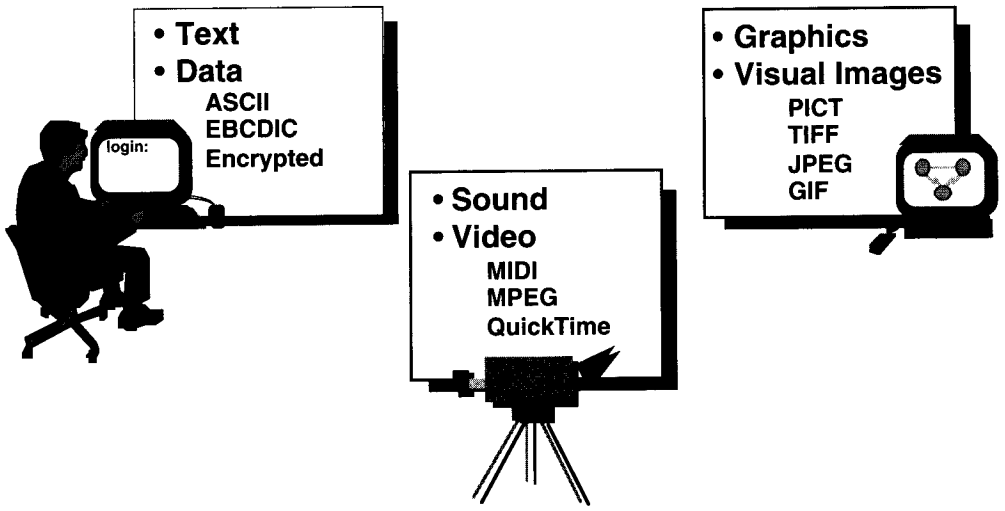

• **Provides code formatting and conversion for applications**

6

The presentation layer provides code formatting and conversion. Code formatting ensures that applications have meaningful information to process. If necessary, the presentation layer translates between multiple data representation formats.

The presentation layer concerns itself not only with the format and representation of actual user data, but also with data structure used by programs; therefore, the presentation layer negotiates data transfer syntax for the application layer.

For example, the presentation layer is responsible for syntax conversion between systems that have differing text and data character representations, such as EBCDIC and ASCII.

Presentation-layer functions also include data encryption. Processes and codes convert data so that the data can be transmitted with its information content protected from unauthorized receivers. Other routines compress text or convert graphic images into bit streams for transmission across a network.

Other Layer 6 standards guide graphic and visual image presentation. PICT is a picture format used to transfer QuickDraw graphics between Macintosh or PowerPC programs. Tagged Image File Format (TIFF) is a standard graphics format for high-resolution, bit-mapped images. JPEG standards come from the Joint Photographic Experts Group.

For sound and movies, presentation layer standards include Musical Instrument Digital Interface (MIDI) for digitized music; also, there is growing acceptance of the Motion Picture Experts Group's (MPEG) standard for compression and coding of motion video for CDs, digital storage, and bit rates up to 1.5 Mbps. QuickTime handles audio and video for Macintosh or PowerPC programs.

Copyright © 1996, Cisco Systems, Inc.

Session Layer

- Network File System (NFS)
- Structured Query Language (SQL)
- Remote-Procedure Call (RPC)
- X Window System
- AppleTalk Session Protocol (ASP)
- DNA Session Control Protocol (SCP)

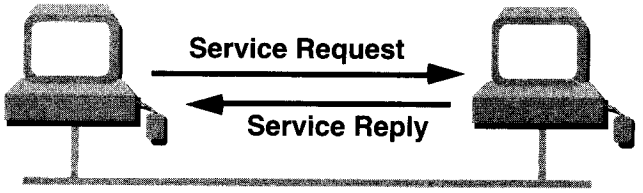

- Coordinates applications as they interact on different hosts

7

The session layer establishes, manages, and terminates sessions between applications. Essentially, the session layer coordinates service requests and responses that occur when applications communicate between different hosts.

Following are examples of session-layer protocols and interfaces:

- Network File System (NFS)—Distributed files system developed by Sun Microsystems to allow transparent access to remote network-based resources; used with TCP/IP and UNIX workstations.
- Structured Query Language (SQL)—Database language developed by IBM to give users an easier way to specify their information needs on local and remote systems.
- Remote procedure call (RPC)—General redirection mechanism for distributed service environments. RPC procedures are built on clients, then executed on servers.
- X Window System—Popular protocol that permits intelligent terminals to communicate with remote UNIX computers as if they were directly attached monitors.
- AppleTalk Session Protocol (ASP)—Establishes and maintains sessions between an AppleTalk client and a server.
- Digital Network Architecture Session Control Protocol (DNA SCP)—DECnet session-layer protocol.

Copyright © 1996, Cisco Systems, Inc.

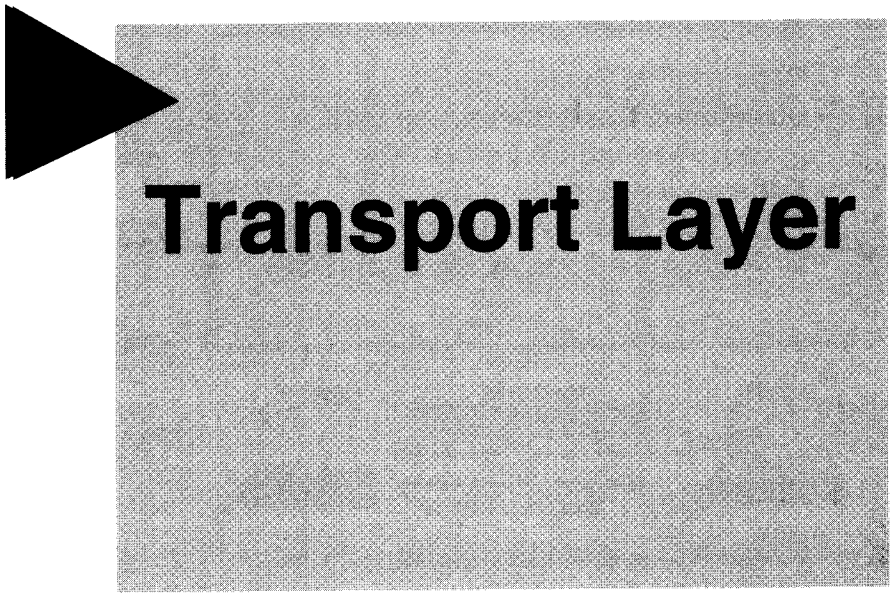

8

Transport Layer

Copyright © 1996, Cisco Systems, Inc.

Transport Layer Overview

- **Segments upper-layer applications**
- **Establishes an end-to-end connection**
- **Sends segments from one end host to another**
- **Optionally, ensures data reliability**

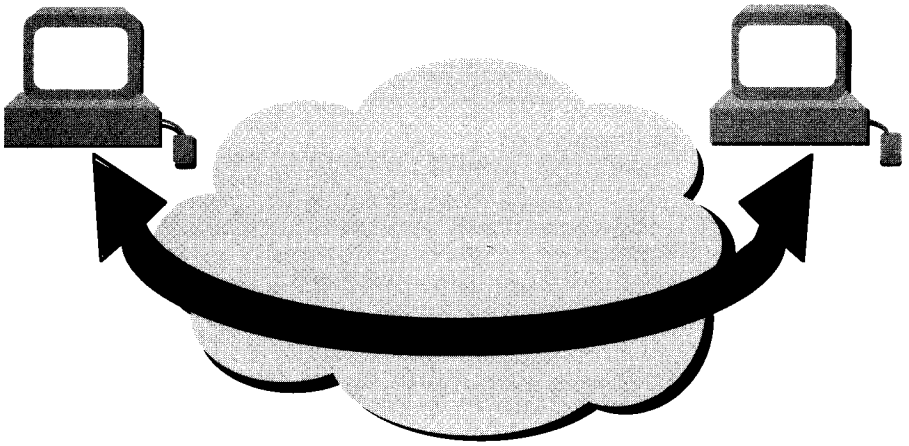

9

Transport services allow users to segment and reassemble several upper-layer applications onto the same transport-layer data stream.

This transport-layer data stream provides end-to-end transport services. It constitutes a logical connection between the endpoints of the internetwork: the originating or sender host and the destination or receiving host.

As the transport layer sends its segments, it can also ensure data integrity. One method provides flow control. Flow control avoids the problem of a host at one side of the connection overflowing the buffers in the host at the other side. Overflows can cause lost data.

Transport services also allow users to request reliable data transport between communicating end systems. Reliable transport uses a connection-oriented relationship between the communicating end systems to accomplish the following:

- Ensure that segments delivered will be acknowledged back to the sender
- Provide for retransmission of any segments that are not acknowledged
- Put segments back into their correct sequence at the destination
- Provide congestion avoidance and control

A more detailed discussion of reliable transport occurs later in this chapter.

Copyright © 1996, Cisco Systems, Inc.

Segment Upper-Layer Applications

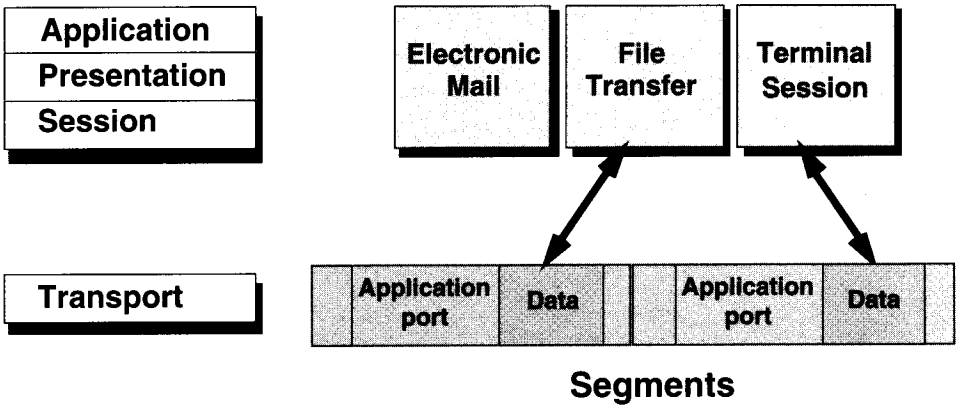

- **Transport segments share traffic stream**

10

One reason for different layers in the OSI reference model is to allow multiple applications to share a transport connection.

Transport functionality is accomplished segment by segment. Each segment is autonomous. Different applications can send successive segments on a first-come, first-served basis. These segments can be intended for the same destination host or many different destination hosts.

For example, several applications from a source host can communicate with corresponding applications on the same destination host, or several applications on an originating host may communicate with corresponding applications on many different destination hosts.

Software in the source machine must set the necessary port number for each software application before transmission. When sending a message, the source computer includes extra bits that encode the message type, originating program, and protocols used.

Then each software application that sends a data stream segment uses the same previously defined port number.

When the destination computer receives the data stream, it can separate and rejoin each application's segments, allowing the transport layer to pass the data up to its destination peer application.

Copyright © 1996, Cisco Systems, Inc.

Establishes Connection

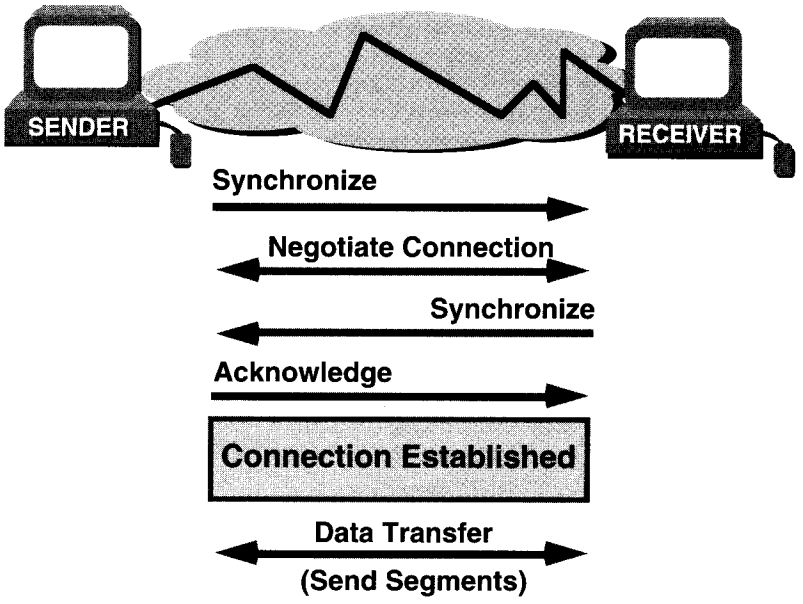

11

To use the reliable transport services, one user of the transport layer must establish a connection-oriented session with its peer system.

For data transfer to begin, both the sending and receiving application programs inform their respective operating systems that a connection will be initiated. In concept, one machine places a call that must be accepted by the other. Protocol software modules in the two operating systems communicate by sending messages across the network to verify that the transfer is authorized and that both sides are ready.

After all synchronization has occurred, a connection is said to be established, and the transfer of information begins. During transfer, the two machines continue to communicate with their protocol software to verify that data is received correctly.

The graphic depicts a typical connection between sending and receiving systems. The first handshake segment requests synchronization. The second and third segments acknowledge the initial synchronization request, as well as synchronize connection parameters in the opposite direction. The final handshake segment is an acknowledgment used to inform the destination that both sides agree that a connection has been established. Once the connection has been established, data transfer begins.

Copyright © 1996, Cisco Systems, Inc.

Sends Segments with Flow Control

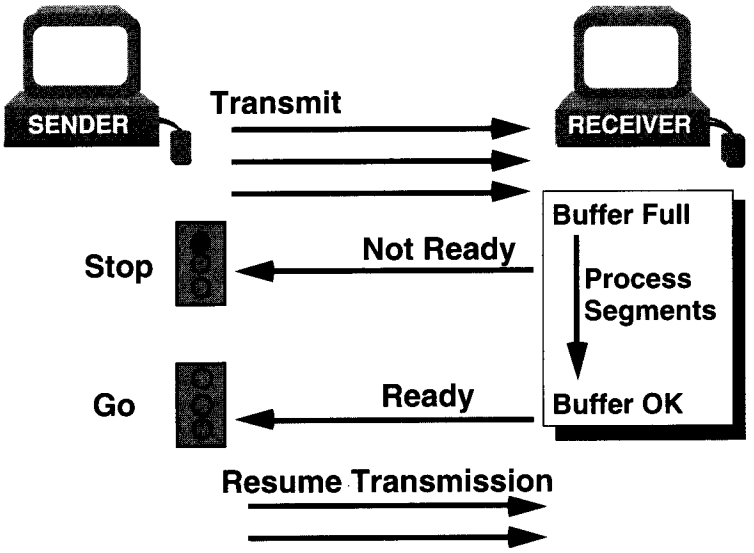

12

Once data transfer is in progress, congestion can arise for two different reasons. First, a high-speed computer might be able to generate traffic faster than a network can transfer it. Second, if many computers simultaneously need to send datagrams through a single gateway or to a single destination, that gateway or destination can experience congestion, even though no single source caused the problem.

When datagrams arrive too quickly for a host or gateway to process, they are stored in memory temporarily. If the datagrams are part of a small burst, this buffering solves the problem. If the traffic continues, the host or gateway eventually exhausts its memory and must discard additional datagrams that arrive.

Instead of allowing data to be lost, the transport function can issue a "not ready" indicator to the sender. Acting like a stop sign, this indicator signals the sender to stop sending segment traffic to its peer. When the peer receiver can handle additional segments, the receiver sends a "ready" transport indicator, which is like a go signal. When it receives this indicator, the sender can resume segment transmission.

Copyright © 1996, Cisco Systems, Inc.

Reliability with Windowing

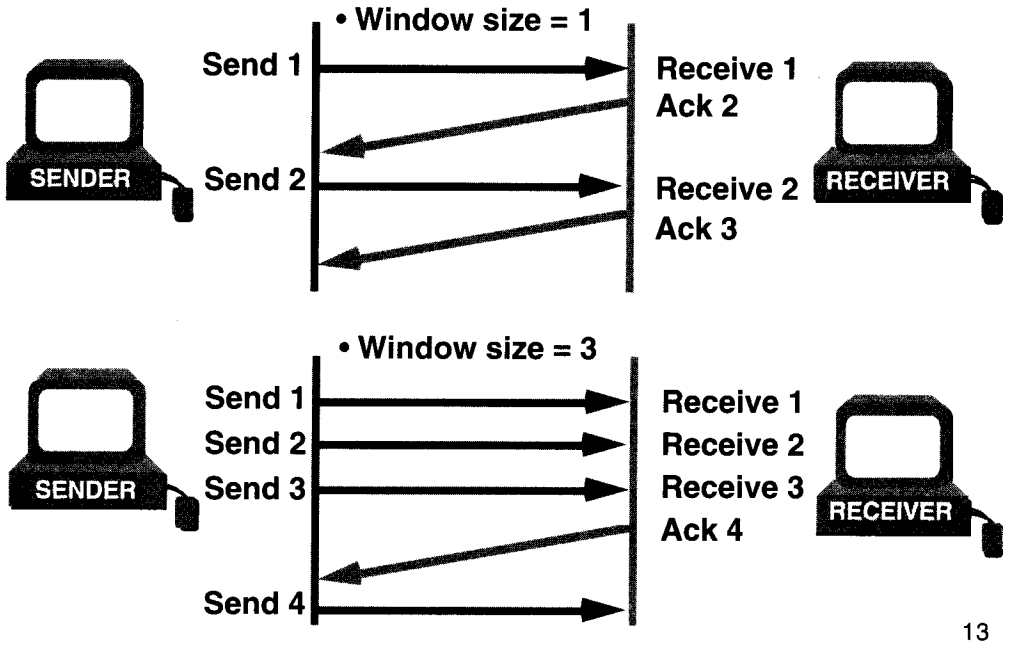

13

In the most basic form of reliable connection-oriented data transfer, data segments must be delivered to the recipient in the same sequence that they were transmitted. The protocol in question fails if any data segments are lost, damaged, duplicated, or received in a different order. The basic solution is to have a recipient acknowledge the receipt of every data segment.

If the sender has to wait for an acknowledgment after sending each segment, throughput will be low. Because time is available after the sender finishes transmitting the data segment and before the sender finishes processing any received acknowledgment, the interval is used for transmitting more data. The number of data segments the sender is allowed to have outstanding—without yet receiving an acknowledgment—is known as the window.

Windowing is a method to control the amount of information transferred end-to-end. Some protocols measure information in terms of the number of packets; TCP/IP measures information in terms of the number of bytes.

In the graphic's examples, a window size of 1 is shown, followed by a window size of 3. With a window size of 1, the sender waits for an acknowledgment for every data segment transmitted. With a window size of 3, the sender can transmit three data segments before expecting an acknowledgment.

Windowing is an end-to-end facility between sender and receiver. In the graphic's example, sender and receiver are workstations. Unlike in this simplified graphic, there is a high probability that acknowledgments and packets will intermix as they communicate across the network through routers. In this example, routers do not intervene in the windowing function between these workstations.

Copyright © 1996, Cisco Systems, Inc.

An Acknowledgment Technique

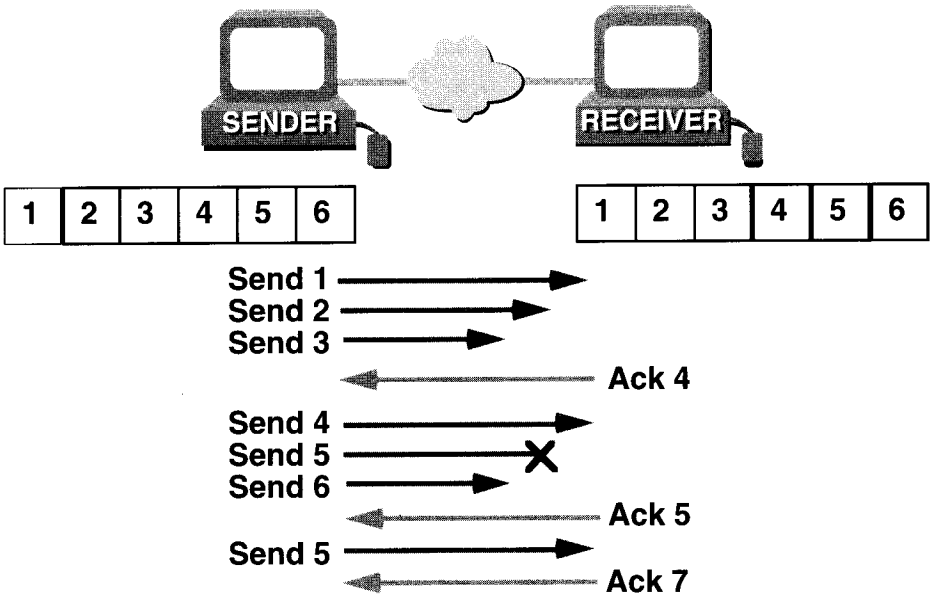

14

Reliable delivery guarantees that a stream of data sent from one machine will be delivered through a functioning data link to another machine without duplication or data loss. Positive acknowledgment with retransmission is one technique that guarantees reliable delivery of data streams. Positive acknowledgment requires a recipient to communicate with the source, sending back an acknowledgment message when it receives data. The sender keeps a record of each segment it sends and waits for an acknowledgment before sending the next segment.

The sender also starts a timer when it sends a segment, and it retransmits a segment if the timer expires before an acknowledgment arrives.

The graphic shows the sender transmitting segment 1, 2, and 3. The receiver acknowledges receipt of the segments by requesting segment number 4. The sender, upon receiving the acknowledgment, sends segments 4, 5, and 6. If segment number 5 does not arrive at the destination, the receiver acknowledges with a request to resend segment number 5. The sender resends segment number 5 and must receive an acknowledgment to continue with the transmission of segment number 7.

Copyright © 1996, Cisco Systems, Inc.

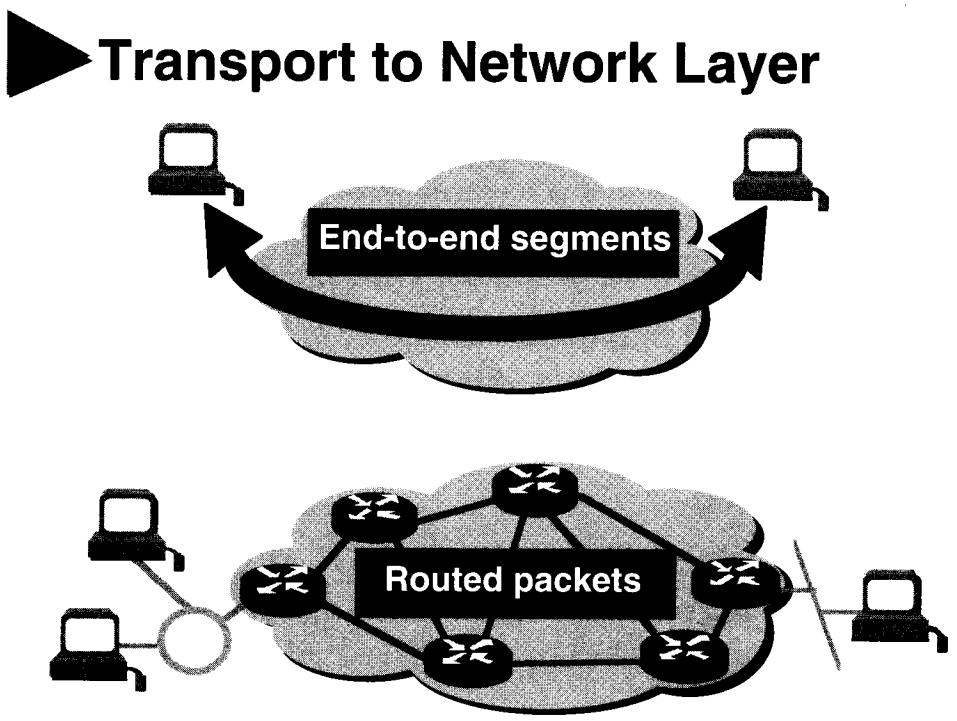

In this chapter, you have learned the upper-layer aspects of network functioning. As each of the upper levels performs its own functions, it depends on lower-layer services as needed. All four upper layers—application (Layer 7), presentation (Layer 6), session (Layer 5), and transport (Layer 4)—can encapsulate data in end-to-end segments.

The transport layer assumes it can use the network as a given cloud as segments cross from sender source to receiver destination.

If we open up the functions inside the cloud, we reveal issues such as "Which of several paths is best for a given route?" We see the role that routers perform in this process, and we see the segments of Layer 4 transport further encapsulated into packets. These issues constitute the focus of the final chapter in this module.

Copyright © 1996, Cisco Systems, Inc.

Summary

The ISO/OSI reference model describes network applications

Presentation layer formats and converts network application data to represent text, graphics, images, video, and audio

Session-layer functions coordinate communication interactions between applications

Reliable transport-layer functions include

- **Multiplexing**
- **Connection synchronization**
- **Flow control**
- **Error recovery**
- **Reliability through windowing**

Copyright © 1996, Cisco Systems, Inc.

Exercise: Applications and Upper Layers

Problem 1

Objective: Name and describe computer, network, and internetwork applications.

Early in this chapter, the presentation distinguished between various application types. Write three examples of network applications and three examples of internetworking applications; include a brief description of what each application does.

A) Network Applications:

B) Internetworking Applications:

Problem 2

Objective: Describe the OSI presentation functions and identify common standards.

Write a description of the OSI presentation-layer functions. Then write the names of at least two presentation-layer standards or types.

Problem 3

Objective: Describe the OSI session functions and identify common standards.

Write a description of the OSI session-layer functions. Then write the names of at least two session-layer standards.

Copyright © 1996, Cisco Systems, Inc.

Problem 4

Objective: Describe the OSI transport functions for end-to-end network services.

The transport function in the ISO/OSI reference model extends services across some portion of the topology shown below. Circle the letter of the brace that indicates the correct extent of transport services.

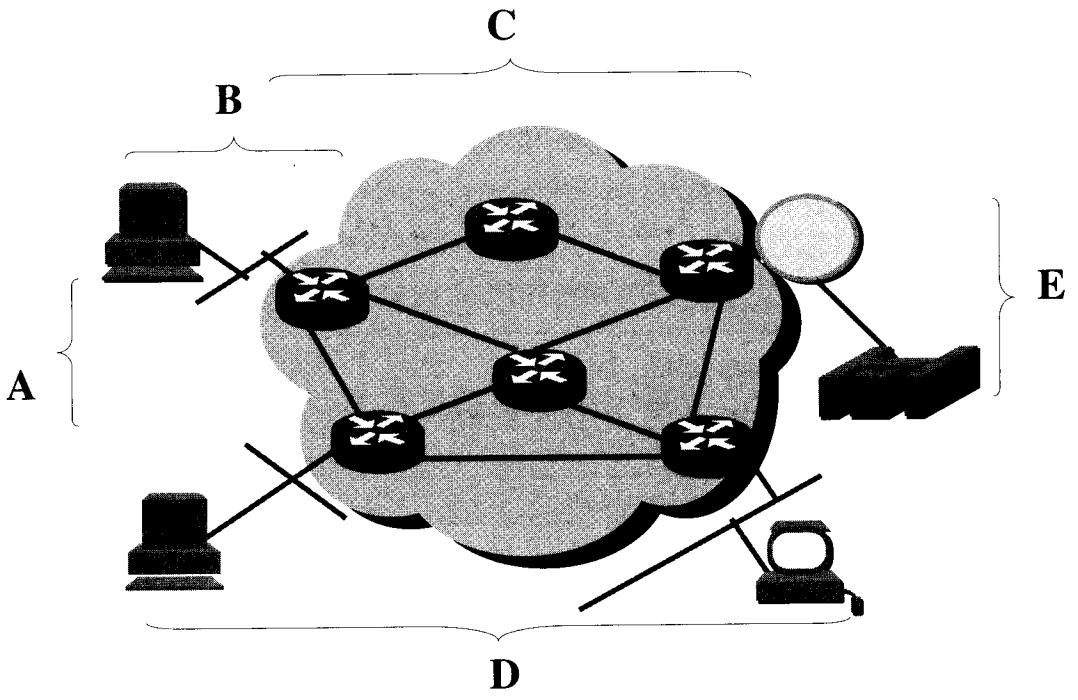

Problem 5

Objective: Identify common processes for establishing connections, flow control, and windowing.

A) Telnet and FTP can share a Layer 4 connection. What process allows this?

_____

_____

B) Data can be lost if a host on one side is allowed to send so many segments that the other-side host overflows its buffers. Write the name of the generic process used to avoid this problem.

_____

_____

C) TCP enables reliability by using windowing. List two mechanisms that work with windowing to help ensure delivery of segments without duplication and data loss.

_____

_____

Copyright © 1996, Cisco Systems, Inc.

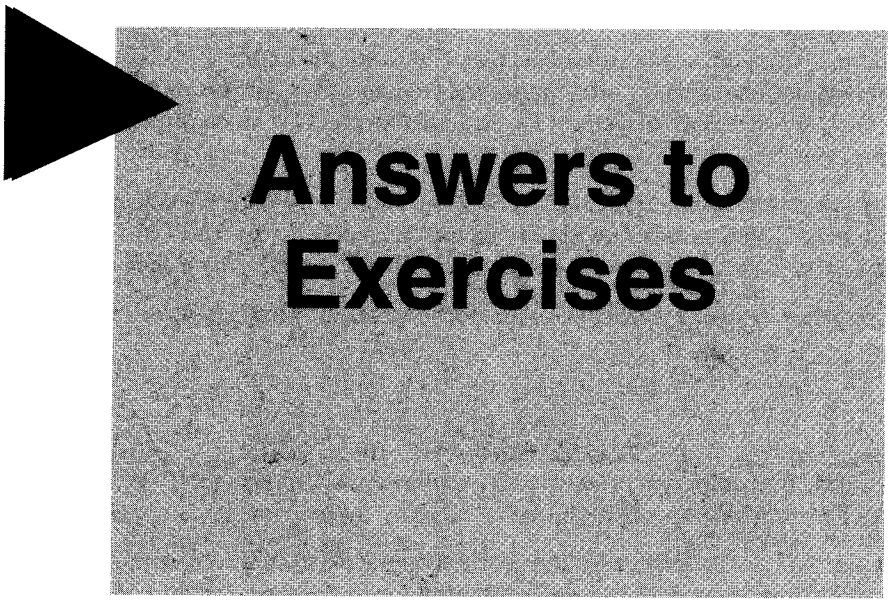

19

Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

Exercise: Applications and Upper Layers

Problem 1

Here are some of the many correct answers for problem 1:

A) Network Applications: Eudora, Quickmail, and other e-mail applications; the mail gateway Simple Mail Transfer Protocol (SMTP); FTP; Telnet; the numerous asynchronous applications such as Procomm Plus, Versaterm, and Crosstalk; Novell's NetWare Core Protocol; ping and trace; Simple Network Management Protocol (SNMP) and Common Management Information Protocol (CMIP) for network management.

B) Internetworking Applications: electronic data interchange (EDI); World Wide Web (WWW) with browsers such as Mosaic and Netscape; e-mail gateways such as the X.400 standard or SMTP; special-interest bulletin boards; financial transaction services; Gopher, Fetch, and Wide Area Information Server (WAIS); video, voice, and data conferencing applications. Also others.

Although the TCP/IP stack merges the top three layers into a single layer, the TCP/IP protocols appear in the list as possible answers for instructional consistency.

Problem 2

HOW INFORMATION IS PRESENTED TO THE USER.

The presentation layer provides code formatting and conversion for applications. Although the TCP/IP stack merges the top three layers into a single layer, the TCP/IP protocols appear as follows as possible answers for instructional consistency:

Standards for text and data include ASCII, EBCDIC, and all types of encryption and data compression. For graphics and visual images, standards include PICT, TIFF, JPEG, and GIFF. Sound and movies use standards including MIDI, MPEG, QuickTime, and HTML.

Problem 3

The session layer coordinates applications as they interact on different hosts.

Again, the TCP/IP protocols appear as follows as possible session-layer answers for instructional consistency:

- Standards include NFS, SQL-based front-end/back-end databases, RPC, the X Window System, Apple's ASP, and Digital's SCP.
- You can also allow NetBIOS or NETBUI as answers here, but avoid spending time on too much tangent discussion.

Problem 4 Answers

D) The transport function extends from end to end. A source host sends segments to a destination host across a given network cloud.

Problem 5 Answers

A) Application segmenting—Port (or in UNIX, socket) numbers enable segments from several upper-layer applications to share the same end-to-end connection.

B) Flow control—A segment message to the sender stops transmission until the receiver's buffers are no longer in danger of filling and then overflowing.

C) Acknowledgment/negative acknowledgment and sequence numbering with retransmission all work to provide reliability.

Copyright © 1996, Cisco Systems, Inc.

CHAPTER 3

Physical and Data Link Layers

Copyright © 1996, Cisco Systems, Inc.

Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

Identify and describe the data link sublayers and their functions

Explain the use of MAC addresses

Describe the topology and functionality of LANs

Differentiate between LAN and WAN protocols

Describe the characteristics of WAN-based protocols

2

This chapter discusses the physical and data link layers of the OSI reference model. It also discusses the operation of the three most commonly used LAN topologies: Ethernet, Token Ring, and FDDI. Following the LAN discussion, common WAN technologies are explained.

The previous chapter presented the top four layers of the OSI reference model. This chapter presents the bottom two layers. The remaining layer, the network layer, is discussed in the next chapter.

Sections:

- Physical and Data Link Layers
- Common LAN Technologies
- Common WAN Technologies
- Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

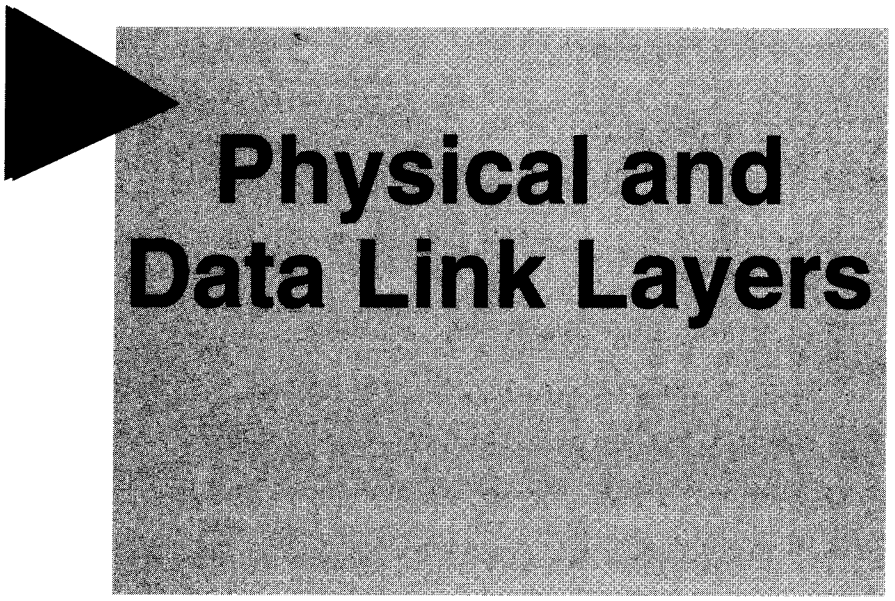

3

Physical and Data Link Layers

Copyright © 1996, Cisco Systems, Inc.

Physical and Data-Link Standards

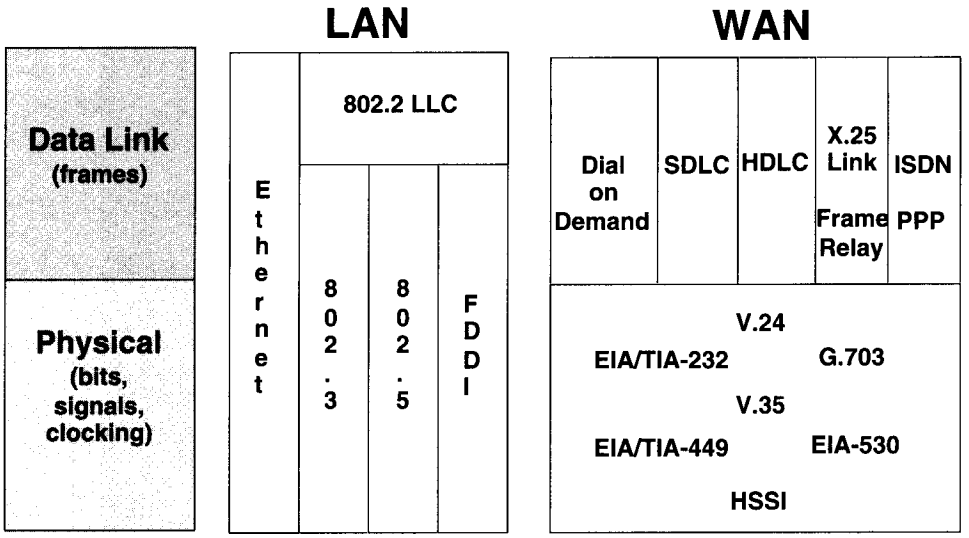

- **Separate physical and data link layers for LAN and WAN**

4

The data link layer provides data transport across a physical link. To do so, the data link layer handles physical addressing, network topology, line discipline, error notification, orderly delivery of frames, and optional flow control.

The physical layer specifies the electrical, mechanical, procedural, and functional requirements for activating, maintaining, and deactivating the physical link between end systems. The physical layer specifies characteristics such as voltage levels, data rates, maximum transmission distances, and physical connectors.

These requirements and characteristics are codified into standards. For example, EIA/TIA-232 standardizes a physical connection to voice-grade access.

You can best understand physical and data link layers by considering WAN and LAN protocols separately. As the graphic shows, certain layer standards are used with LAN links, and certain other layer standards are used by WAN links.

For example, a dial-on-demand protocol places a WAN call based on protocol traffic defined as "interesting." However, the dial-on-demand protocol, opening a WAN call for WAN bandwidth, has no direct interface with the LAN physical or data-link operations where that traffic might have originated.

Instead of LAN operations and framing, dial-on-demand routing (DDR) opens and controls a "pipe" of bandwidth on a WAN physical interfaces. For instance, DDR might place a call over the physical interface of V.35, which is recommended for high-speed WAN access.

Copyright © 1996, Cisco Systems, Inc.

LAN Data Link Sublayers

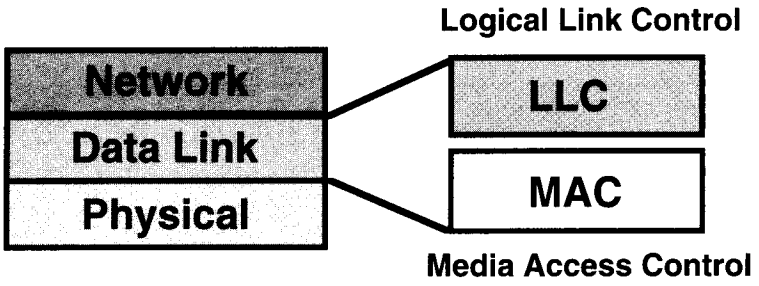

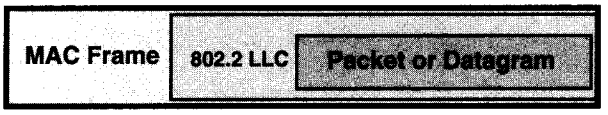

- **LLC refers upward to higher-layer *software* functions**
- **MAC refers downward to lower-layer *hardware* functions**

5

LAN protocols occupy the bottom two layers of the OSI reference model: the physical layer and data link layer. The Institute of Electrical and Electronic Engineers (IEEE) 802 committee subdivided the data link layer into two sublayers: the logical link control (LLC) sublayer and the media access control (MAC) sublayer. The graphic illustrates the layers. We will cover each sublayer in sequence.

The LLC sublayer provides for environments that need connectionless or connection-oriented services at the data link layer.

The MAC sublayer provides access to the LAN medium in an orderly manner.

Copyright © 1996, Cisco Systems, Inc.

LLC Sublayer Functions

- **Enable upper layers to gain independence over LAN media access**
- **Allow service access points (SAPs) from interface sublayers to upper-layer functions**
- **Provide optional connection, flow control, and sequencing services**

6

The LLC sublayer rests on top of the other 802 protocols to provide interface flexibility. Upper-layer protocols, for example IP at Layer 3, can operate autonomously without regard for the specific type of LAN media. This independence occurs because, unlike the MAC sublayer, LLC is not limited to a specific 802 MAC protocol. Instead, the LLC sublayer can depend on lower layers to provide access to the media.

From the perspective of these lower MAC sublayers, the SAP process provides a convenient interface to the upper OSI layers. These SAP entries simplify access to the shared channel up to the specified upper-layer service identified by LLC SAP entities.

LLC sublayer options include support for connections between applications running on the LAN, flow control to the upper layer by means of ready/not ready codes, and sequence control bits.

Copyright © 1996, Cisco Systems, Inc.

Physical and Logical Addressing

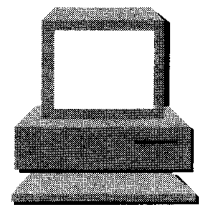

0000.0c12.3456

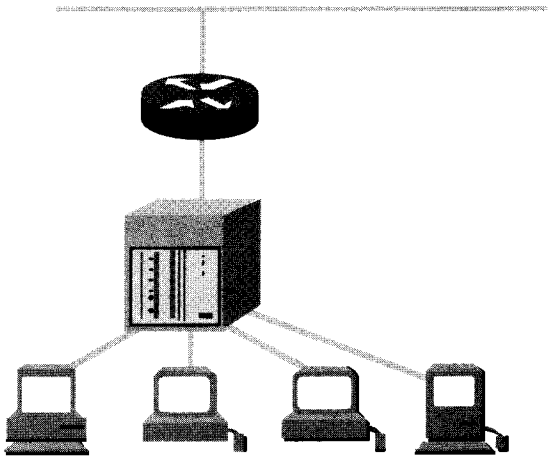

7

Locating computer systems on an internetwork is an essential component of any network system. There are various addressing schemes used for this purpose, depending on the protocol family being used. In other words, AppleTalk addressing is different from TCP/IP addressing, which in turn is different from OSI addressing, and so on.

Two important types of addresses are link-layer addresses and network-layer addresses. Link-layer addresses (also called physical or hardware addresses) are typically unique for each network connection. In fact, for most local-area networks (LANs), link-layer addresses are resident in the interface circuitry and are assigned by the organization that defined the protocol standard represented by the interface. Because most computer systems have one physical network connection, they have only a single link-layer address. Routers and other systems connected to multiple physical networks can have multiple link-layer addresses. As their name implies, link-layer addresses exist at Layer 2 of the OSI reference model.

Network-layer addresses (also called virtual or logical addresses) exist at Layer 3 of the OSI reference model. Unlike link-layer addresses, which usually exist within a flat address space, network-layer addresses are usually hierarchical. In other words, they are like mail addresses, which describe a person's location by providing a country, a state, a zip code, a city, a street, street address, and finally, a name. One good example of a flat address space is the U.S. Social Security numbering system, where each person has a single, unique social security number.

Copyright © 1996, Cisco Systems, Inc.

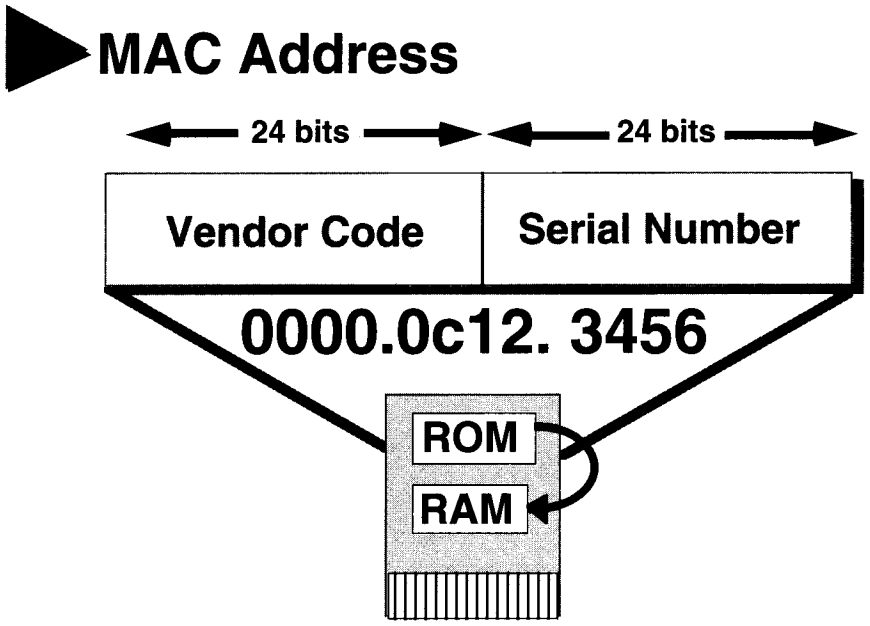

For multiple stations to share the same medium and still uniquely identify each other, the MAC sublayer defines a hardware or data-link address called the MAC address. The MAC address is unique for each LAN interface.

On most LAN-interface cards, the MAC address is burned into ROM—hence the term burned-in address (BIA). When the network interface card initializes, this address is copied into RAM.

The MAC address is a 48-bit address expressed as 12 hexadecimal digits. The first 6 hexadecimal digits of a MAC address contain a manufacturer identification (vendor code) also known as the Organizational Unique Identifier (OUI). To ensure vendor uniqueness, the IEEE administers OUIs. The last 6 hexadecimal digits are administered by each vendor and often represent the interface serial number.

Copyright © 1996, Cisco Systems, Inc.

Finding the MAC Address

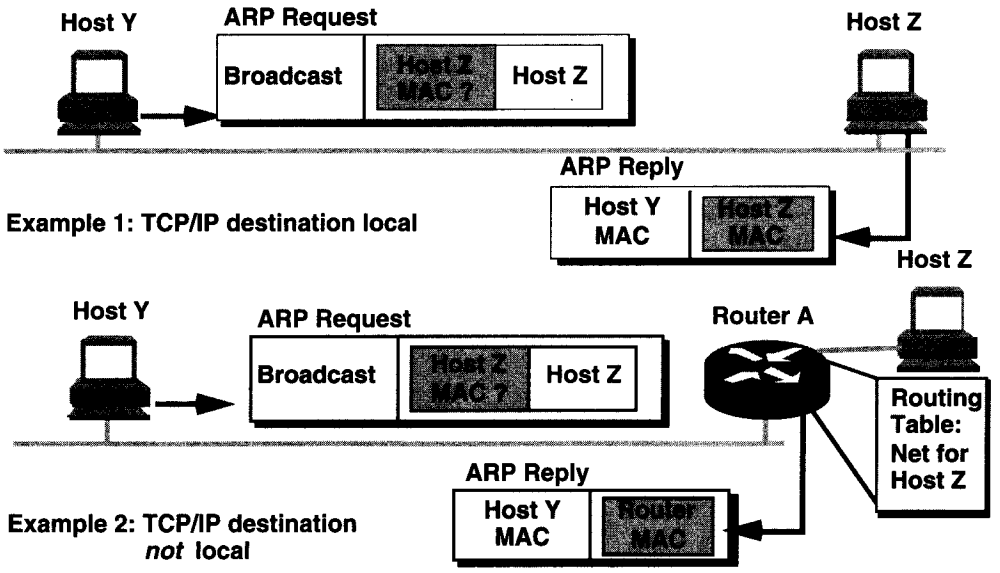

- **An example: TCP/IP Address Resolution Protocol (ARP)**
- **ARP finds the MAC address for a data-link connection**

9

Before a frame is exchanged with a directly connected device, the sending device needs to have a MAC address it can use as a destination address. One way to discover a device's MAC address is to use an address resolution protocol. The graphic illustrates two ways in which a TCP/IP example, ARP, is used to discover a MAC address.

In the first example, host Y and host Z are on the same LAN. Host Y broadcasts an ARP request to the LAN looking for host Z. Because host Y has sent out a broadcast, all devices including host Z will process the request; however, only host Z will respond with its MAC address. Host Y receives host Z's reply and saves the MAC address in local memory, often called an ARP cache. The next time host Y needs to directly communicate with host Z, it recalls host Z's stored MAC address.

In the second example, host Y and host Z are on different LANs, but can access each other through router A. When host Y broadcasts its ARP request, router A determines that host Z cannot recognize the request because router A knows that host Z is on a different LAN. Because router A further determines that any packets for host Z must be relayed, router A provides its own MAC address as a proxy reply to the ARP request. Host Y receives the router's response and saves the MAC address in its ARP cache memory. The next time host Y needs to communicate with host Z, it recalls the stored MAC address of router A.

Copyright © 1996, Cisco Systems, Inc.

Exercise: MAC Addresses

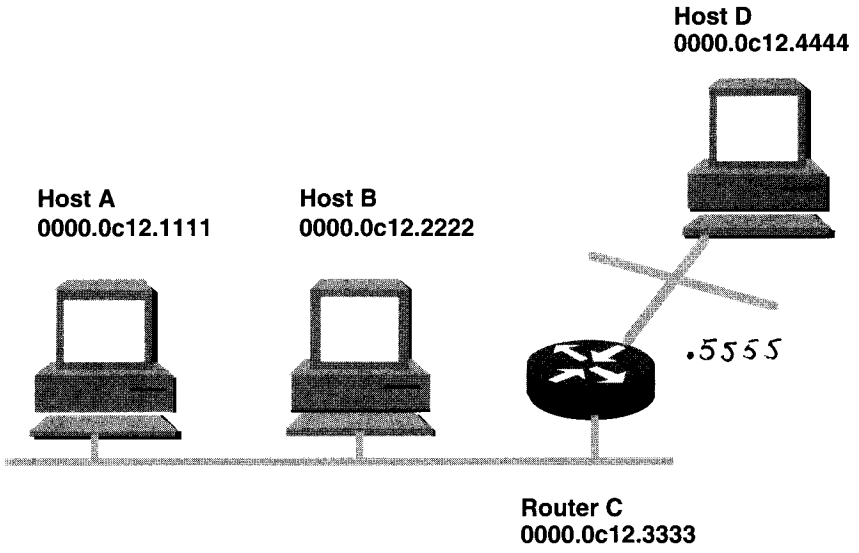

10

Exercise: MAC Addresses

Objective: Explain the use of MAC addresses.

Complete the table to show what address each device would have for another device.

Host A addresses:

- for host B _____
- for host D _____

Host B addresses:

- for host A _____
- for host D _____

Host D addresses:

- for host A _____
- for host B _____

Router C addresses:

- for host A _____
- for host B _____
- for host D _____

Answers 3-41

Copyright © 1996, Cisco Systems, Inc.

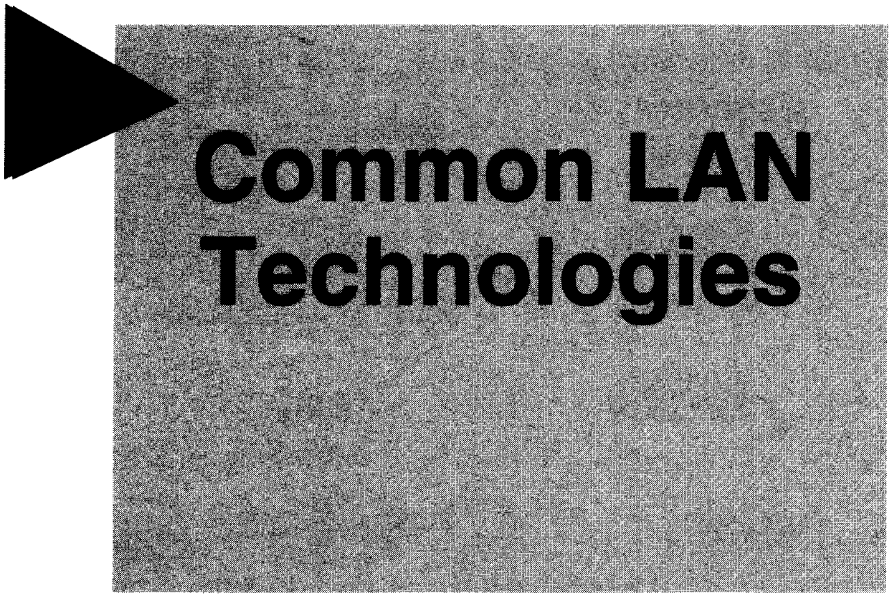

11

Common LAN Technologies

Copyright © 1996, Cisco Systems, Inc.

LAN Technology Overview

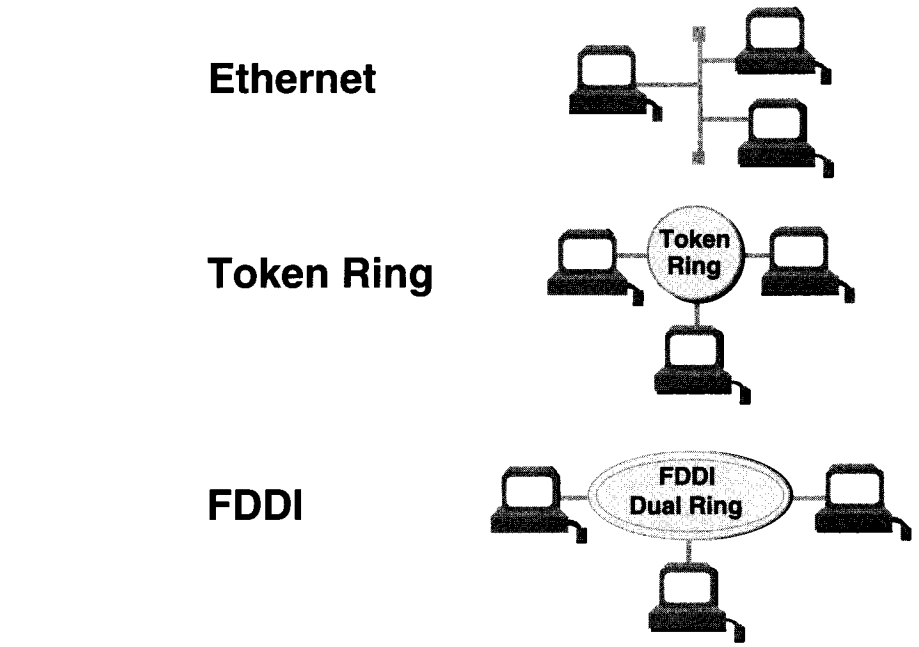

12

You will now learn specific information about the common LAN technologies. The three LAN technologies shown in the graphic account for virtually all deployed LANs:

- Ethernet—The first of the major LAN technologies, it runs the largest number of LANs.
- Token Ring—From IBM, it followed Ethernet and is now widely used in a large number of IBM networks.
- FDDI—Also using tokens, it is now a popular campus LAN.

Pages that follow introduce each technology and describe the physical and data-link details of each.

Copyright © 1996, Cisco Systems, Inc.

Ethernet and IEEE 802.3

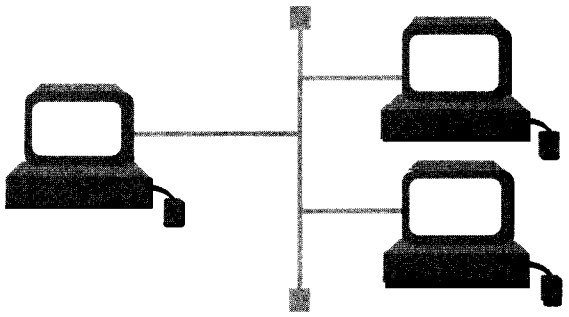

- **Several framing variations exist for this common LAN technology**

13

Xerox performed initial development of Ethernet and was joined by the Digital Equipment Corporation (Digital) and Intel to define the Ethernet 1 specification in 1980. The same group subsequently released the Ethernet II specification in 1984. The Ethernet specification describes a carrier sense multiple access collision detect (CSMA/CD) LAN.

The IEEE 802.3 subcommittee adopted Ethernet as its model for its CSMA/CD LAN specification. As a result, Ethernet II and IEEE 802.3 are identical in the way they use the physical medium.

However, the two specifications differ in their descriptions of the data link layer. These differences do not prohibit manufacturers from developing network interface cards that support the common physical layer, MAC addressing, and software that recognizes the differences between the two logical link control layers.

Copyright © 1996, Cisco Systems, Inc.

Physical Layer: Ethernet/802.3

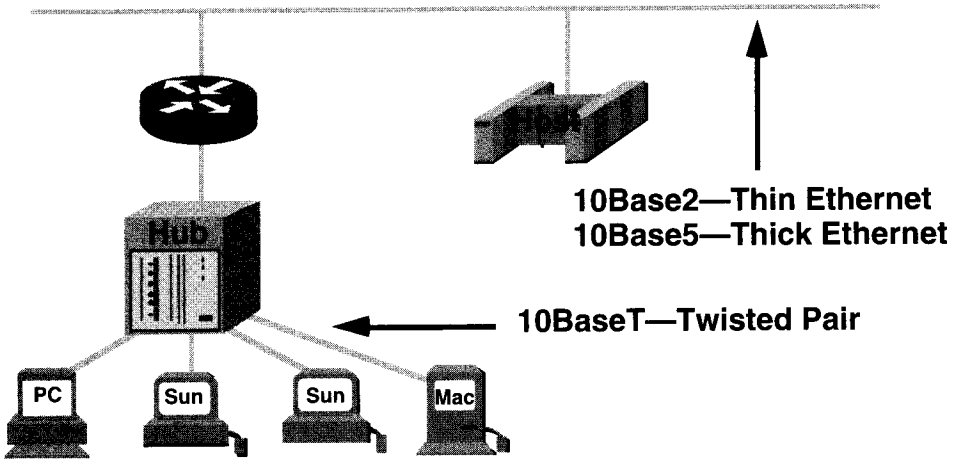

14

The Ethernet and IEEE 802.3 standards define a bus-topology LAN that operates at a baseband signaling rate of 10 Mbps. (An earlier version of Ethernet that operated at 3 Mbps is now obsolete. Newer versions operating at higher speeds are under development.) The graphic illustrates the three defined wiring standards:

- 10Base2—Known as thin Ethernet—allows network segments up to 185 meters on coaxial cable.
- 10Base5—Known as thick Ethernet—allows network segments up to 500 meters on coaxial cable.
- 10BaseT—Carries Ethernet frames on inexpensive twisted-pair wiring.

The 10Base5 and 10Base2 standards provide access for several stations on the same segment. Stations are attached to the segment by a cable that runs from an attachment unit interface (AUI) in the station to a transceiver that is directly attached to the Ethernet coaxial cable. In some interfaces, the AUI and the transceiver are collocated in the station itself, in which case no cable is required.

Because the 10BaseT standard provides access for a single station only, stations attached to an Ethernet LAN by 10BaseT are almost always connected to a hub. In a hub arrangement, the hub is analogous to an Ethernet segment.

Copyright © 1996, Cisco Systems, Inc.

The Ethernet/802.3 Interface

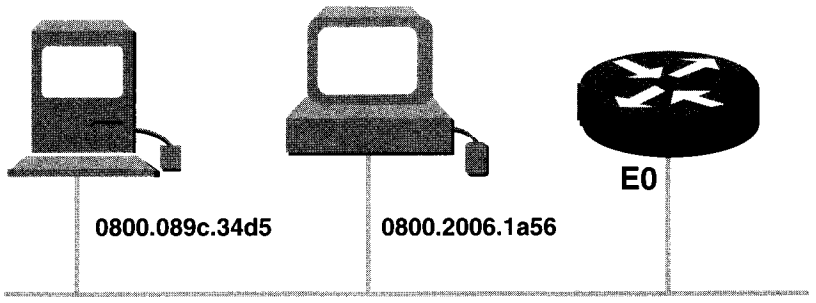

- **Cisco router's data link to Ethernet/802.3 uses an interface named E plus a number (for example, E0)**

15

The Ethernet and 802.3 data links provide data transport across the physical link joining two devices. For example, as this graphic shows, the three devices can be directly attached to each other over the Ethernet LAN.

The Apple Macintosh on the left and the Intel-based PC in the middle show MAC addresses used by the data-link framing. The router on the right also uses MAC addresses for each of its LAN side interfaces. For indicating the 802.3 interface on the router, you will use the Cisco IOS interface type abbreviation E followed by an interface number (for example, 0, as shown in the graphic).

Copyright © 1996, Cisco Systems, Inc.

Ethernet/802.3 Operation

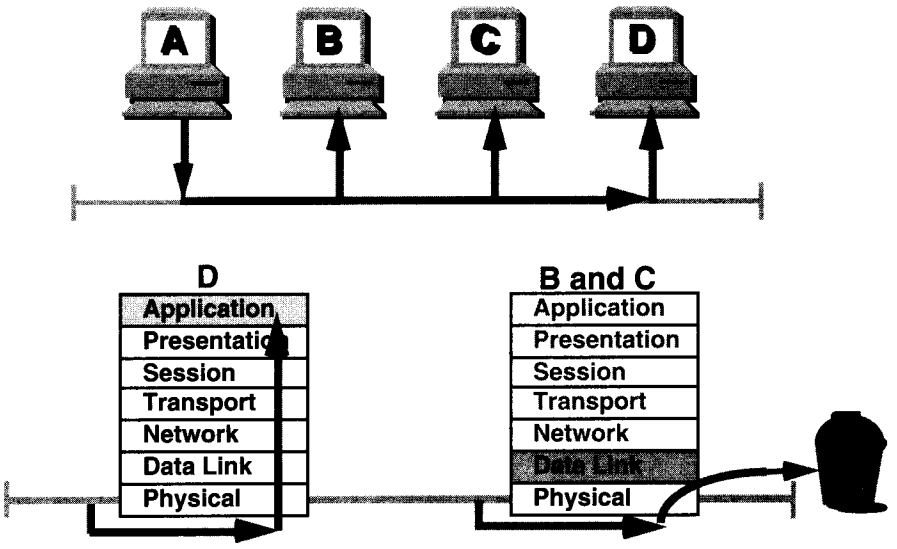

16

In a CSMA/CD network, one node's transmission traverses the entire network and is received and examined by every node. When the signal reaches the end of a segment, terminators absorb it to prevent it from going back onto the segment.

Copyright © 1996, Cisco Systems, Inc.

Ethernet/802.3 Broadcast

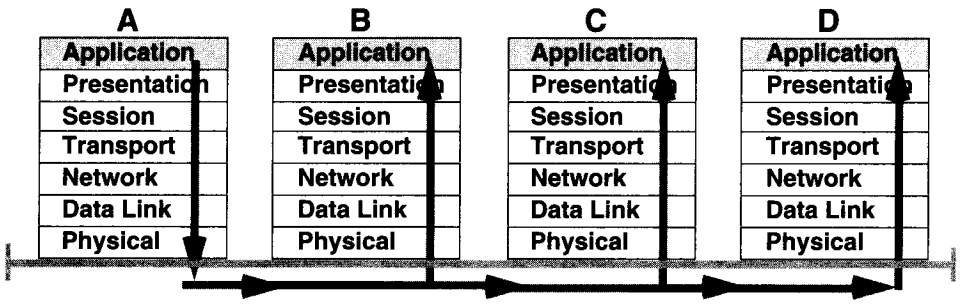

17

Broadcasting is a powerful tool that sends a single frame to many stations at the same time. Broadcasting uses a data-link destination address of all ones (FFFF.FFFF.FFFF in hexadecimal). As the graphic shows, if station A transmits a frame with a destination address of all ones, stations B, C, and D will all receive and pass the frame to their respective upper layers for further processing.

When improperly used, however, broadcasting can seriously impact the performance of stations by interrupting them unnecessarily. For this reason, broadcasts should be used only when the MAC address of the destination is unknown or when the destination is all stations.

A multicast address is a MAC address used to identify a group of destinations and is indicated by the first transmitted bit of the destination address being set to 1. For Ethernet, this bit may appear as the low-order bit (for example, xxxx.xxx1).

Copyright © 1996, Cisco Systems, Inc.

Ethernet Frame Variations

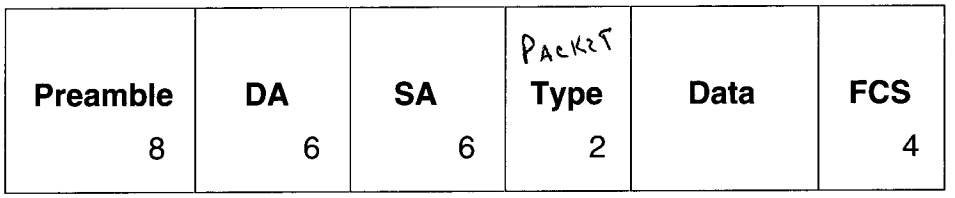

Ethernet Frame

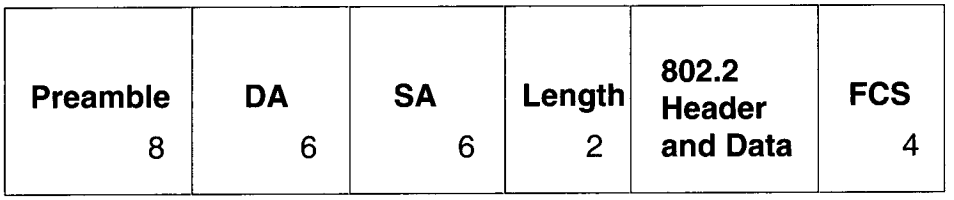

802.3 Frame

18

Both Ethernet and IEEE 802.3 frames begin with an alternating pattern of ones and zeros called a preamble. The preamble tells receiving stations that a frame is coming.

Immediately following the preamble in both Ethernet and IEEE 802.3 LANs are the destination and source physical address fields. Both Ethernet and IEEE 802.3 addresses are six bytes long. Addresses are contained in hardware on the Ethernet and IEEE 802.3 interface cards. The first three bytes are specified by the Ethernet or IEEE 802.3 vendor. The source address is always a unicast (single node) address, while the destination address may be unicast, multicast (group), or broadcast (all nodes).

In Ethernet frames, the two-byte field following the source address is a type field. This field specifies the upper-layer protocol to receive the data after Ethernet processing is complete.

In IEEE 802.3 frames, the two-byte field following the source address is a length field, which indicates the number of bytes of data that follow this field and precede the frame check sequence (FCS) field.

The actual data contained in the frame follows the type/length field. After physical layer and link-layer processing is complete, this data will eventually be sent to an upper-layer protocol.

Following the data field is a four-byte FCS field containing a cyclic redundancy check (CRC) value. The CRC is created by the sending device and recalculated by the receiving device to check for damage that might have occurred to the frame in transit.

Following the length field there is usually an 802.2 header for LLC.

Copyright © 1996, Cisco Systems, Inc.

Ethernet/802.3 Reliability

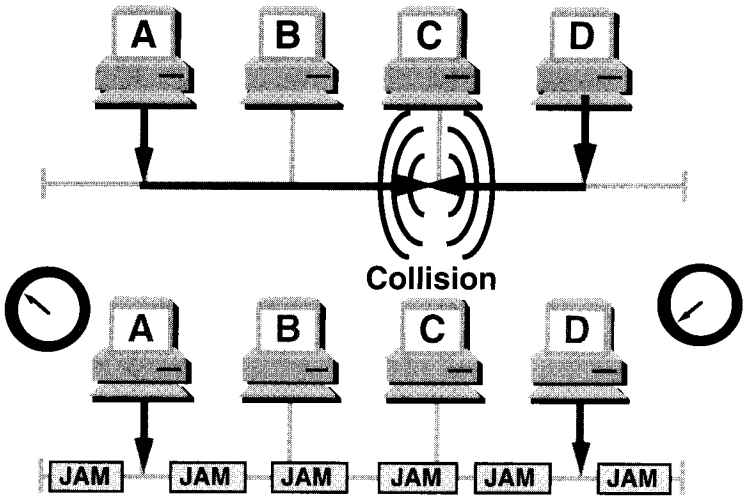

- **Carrier sense multiple access collision detect (CSMA/CD)**

19

CSMA/CD works in the following way: When a station wishes to transmit, it checks the network to determine whether another station is currently transmitting. If the network is not being used, the station proceeds with the transmission. While sending, the station monitors the network to ensure that no other station is transmitting. Two stations might start transmitting at approximately the same time if they determine that the network is available. If two stations send at the same time, a collision occurs, as illustrated in the upper part of the graphic.

When a transmitting node recognizes a collision, it transmits a jam signal that causes the collision to last long enough for all other nodes to recognize it. All transmitting nodes then stop sending frames for a randomly selected time period before attempting to retransmit. If subsequent attempts also result in collisions, the node tries to retransmit up to 15 times before giving up.

The clocks indicate different backoff timers. If the two timers are sufficiently different, one station will succeed the next time. The mean backoff time doubles with each consecutive collision, thereby reducing the chance of collision by the inverse power of 2.

Copyright © 1996, Cisco Systems, Inc.

High-Speed Ethernet Options

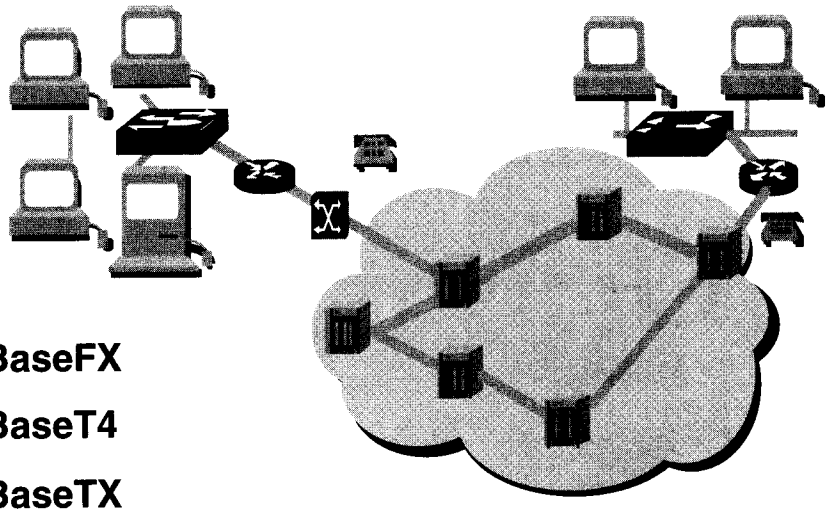

- **100BaseFX**
- **100BaseT4**
- **100BaseTX**
- **100BaseVG AnyLAN**

20

New applications can cause end users to experience delay and other problems such as insufficient bandwidth between end stations. In response to these problems, Ethernet networks are poised to move forward again, with the availability of new, 100-Mbps technologies, such as these:

- 100BaseFX—A 100-Mbps implementation of Ethernet over fiber-optic cable. The MAC layer is compatible with the 802.3 MAC layer.
- 100BaseT4—A 100-Mbps implementation of Ethernet using four-pair Category 3, 4, or 5 cabling. The MAC layer is compatible with the 802.3 MAC layer.
- 100BaseTX—A 100-Mbps implementation of Ethernet over Category 5 and Type 1 cabling. The MAC layer is compatible with the 802.3 MAC layer.
- 100BaseVG AnyLAN—The IEEE specification for 100-Mbps implementation of Ethernet and Token Ring over four-pair UTP. The MAC layer is not compatible with the 802.3 MAC layer.

These high-speed options do not use CSMA/CD like lower-speed Ethernet. Instead, they use 4B/5B signaling.

Copyright © 1996, Cisco Systems, Inc.

Token Ring and IEEE 802.5

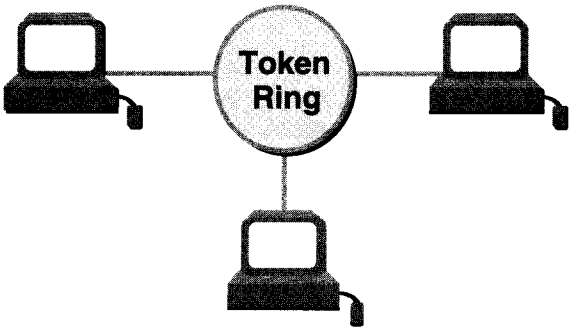

- **IBM's Token Ring is equivalent to IEEE 802.5**

21

Token Ring was developed by IBM in the 1970s. It is still IBM's primary LAN technology and is second only to Ethernet/IEEE 802.3 in popularity. The IEEE 802.5 specification is almost identical to, and completely compatible with, IBM's Token Ring. Both Token Ring specifications are now administered by the IEEE 802.5 committee. The term Token Ring is generally used to refer to both IBM's Token Ring network and IEEE 802.5 networks.

Copyright © 1996, Cisco Systems, Inc.

Physical Layer: Token Ring/802.5

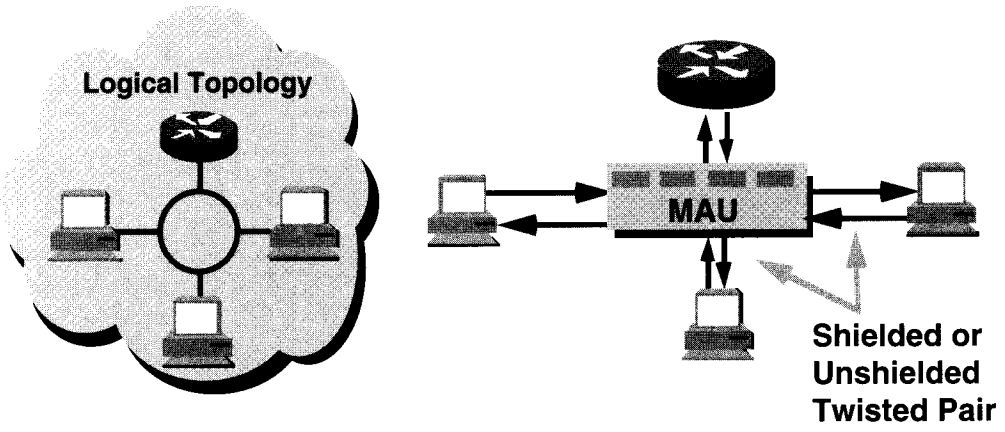

- **Logically a ring, but physically a star configuration to MAU relays**

22

The logical topology of an 802.5 network is a ring in which each station receives signals from its nearest active upstream neighbor (NAUN) and repeats those signals to its downstream neighbor. Physically, however, 802.5 networks are laid out as stars, with each station connecting to a central hub called a multistation access unit (MSAU). This configuration is illustrated in the graphic. The stations connect to the central hub through shielded or unshielded twisted-pair wire.

Typically, an MSAU connects up to eight Token Ring stations. If a Token Ring network has more stations than a MSAU can handle, or if stations are located in different parts of a building—for example, on different floors—MSAUs can be chained together to create an extended ring. When installing an extended ring, you must ensure that the MSAUs themselves are oriented in a ring. Otherwise, the Token Ring will have a break in it and will not operate.

Copyright © 1996, Cisco Systems, Inc.

The Token Ring/802.5 Interface

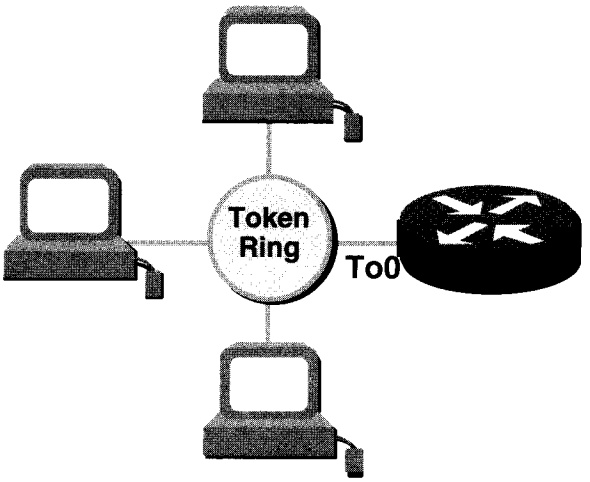

- **Cisco router's data link to Token Ring/802.5 uses interface named *To* plus a number (for example, To0)**

23

The IEEE 802.5 Token Ring protocol parallels IEEE 802.3 by providing MAC sublayer and physical layer services. Token Ring relies on the IEEE 802.2 Logical Link Control (LLC) sublayer and upper-layer protocols for point-to-point services. Token Ring differs considerably from 802.3 in its use of the LAN medium.

All Token Ring stations use MAC addresses, including the router on the right of the graphic. For indicating the 802.5 interface on the router, you will use the Cisco IOS software interface type abbreviation for *token ring (To)* followed by an interface number (for example, *0,* as shown in the graphic).

Copyright © 1996, Cisco Systems, Inc.

Token Ring/802.5 Operation

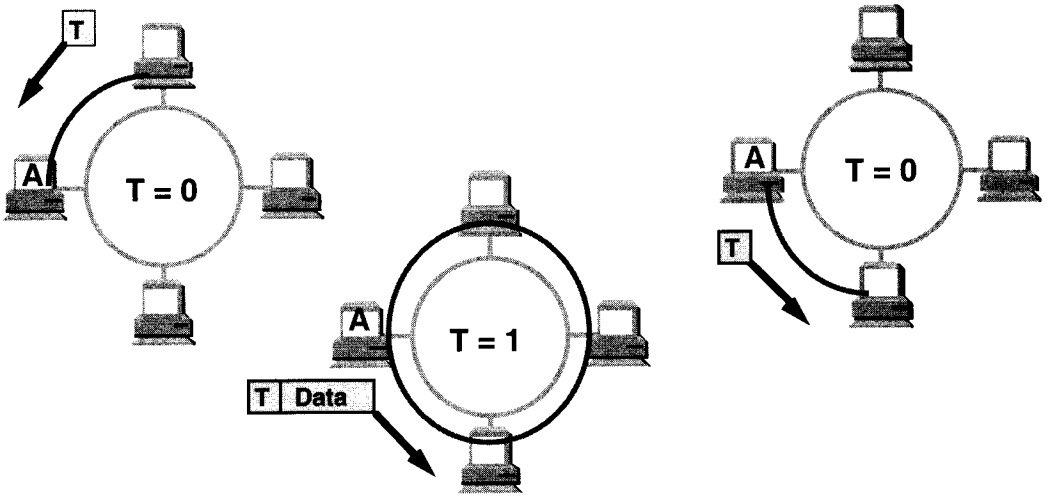

- **Token Ring LANs continuously pass a token or a Token Ring frame**

24

Station access to a Token Ring is deterministic; a station can transmit only when it receives a special frame called a token. Although exceptions can be negotiated, stations are allowed to transmit a single frame when they possess the token. Because no station can dominate the cable as it can in a contention-based access (CSMA/CD) network, administrators can quite accurately determine and plan network performance.

If a station receiving the token has no information to send, it simply passes the token to the next station. If a station possessing the token has information to transmit, it claims the token by altering one bit of the frame, the T bit. The station then appends the information it wishes to transmit and sends the information frame to the next station on the Token Ring.

The information frame circulates the ring until it reaches the destination station, where the frame is copied by the station and tagged as having been copied. The information frame continues around the ring until it returns to the station that originated it, and is removed.

Unless early token release is used on the Token Ring, only one frame can be circling the Token Ring at any one time; other stations wishing to transmit must wait. With early token release, a station that seizes a token can transmit a new token onto the Token Ring after first sending its information frame.

Because frames proceed serially around the ring, and because a station must claim the token before transmitting, collisions are not expected in a Token Ring network.

Broadcasting for source-route bridging is supported in the form of a special mechanism known as explorer packets. These are used to locate a route to a destination through one or more source-route bridges.

Copyright © 1996, Cisco Systems, Inc.

Token Ring/802.5 Media Control

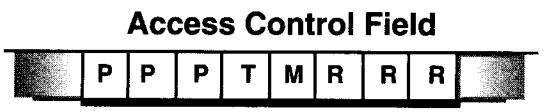

P Priority bits

T Token bit

M Monitor bit

R Reservation bits

- **Fields in a frame determine priority and reservation for sharing media**

25

Token Ring networks use a priority system that permits certain user-designated, high-priority stations to use the network more frequently. Token Ring frames have two fields within the access control field that control priority: the priority field and the reservation field.

Only stations with a priority equal to or higher than the priority of a token can claim that token. After the token is claimed and changed to an information frame, only stations with priority higher than the transmitting station can reserve the token for the next pass around the network. When the next token is generated, it includes the highest priority of the reserving station. Stations that raise a token's priority level must reinstate the previous lower priority level after their transmission is complete.

The graphic illustrates the bits in the access control field that are used to define the current priority and reservation priority.

Copyright © 1996, Cisco Systems, Inc.

Token Ring/802.5 Active Monitor

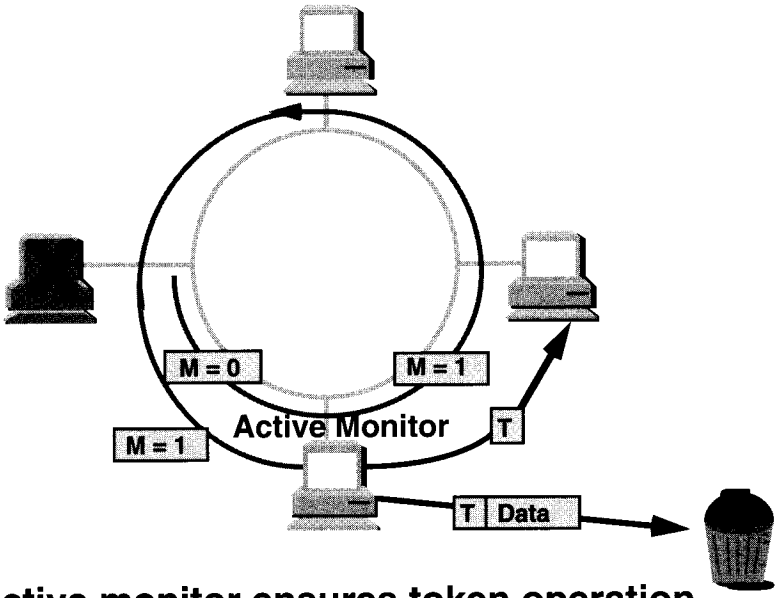

- **Active monitor ensures token operation on the ring for media access**

26

Token Ring networks employ several mechanisms for detecting and compensating for network faults. For example, one station in the Token Ring network is selected to be the active monitor. This station, which can be any station on the network, acts as a centralized source of timing information for other stations and performs a variety of ring maintenance functions.

One ring maintenance function is to remove continuously circulating frames from the ring. When an originating station fails, it is not able to remove its frame from the Token Ring. The leftover frame, which continues to circle the ring, can prevent other stations from transmitting their own frames and can tie up the network. The active monitor can detect such frames, remove them from the ring, and generate a new token.

Copyright © 1996, Cisco Systems, Inc.

Token Ring/802.5 Reliability

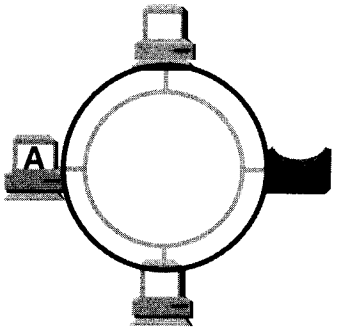

Frame Status Field

	A	C	r	r	A	C	r	r

0	0	Destination not found
0	1	Copied but not acknowledged
1	0	Unable to copy data from frame
1	1	Station found or frame copied to another ring by a bridge

- **Sending station receives status information in a frame**

The IEEE 802.5 specification describes two bits in the frame status field: the A bit, which stands for address (destination MAC address recognized), and the C bit, which stands for copied (the Token Ring frame copied at the destination).

These two bits are used to indicate the status of an outstanding frame. When the Token Ring frame returns to the sender, these bits provide a dependable method for ensuring the sender about the disposition of frames sent out onto the Token Ring.

An originating station generates a frame with the A and C bits turned off (set to zero). Because the originating station always views the returning frame, it can examine these two bits to determine whether they have been modified during their journey around the ring.

Copyright © 1996, Cisco Systems, Inc.

Fiber Distributed Data Interface (FDDI)

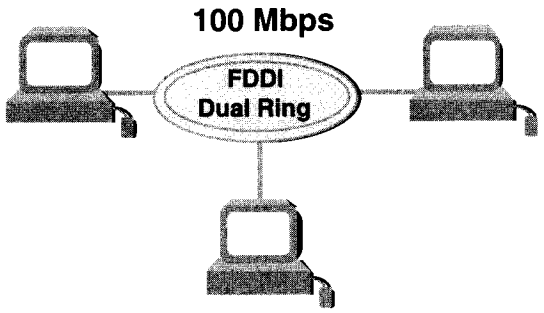

- **Devices on FDDI maintain connectivity on dual counter-rotating rings**

28

FDDI is an American National Standards Institute (ANSI) standard that defines a dual Token Ring LAN operating at 100 Mbps over a fiber-optic medium. The FDDI standards were published in 1987 in the ANSI X3T9.5 standards.

Note ANSI has defined a Twisted-Pair Physical Medium Dependent standard. Based on this standard, Copper Distributed Data Interface (CDDI) provides operation of FDDI but using the more commonly used copper cabling.

Copyright © 1996, Cisco Systems, Inc.

Physical Layer: FDDI

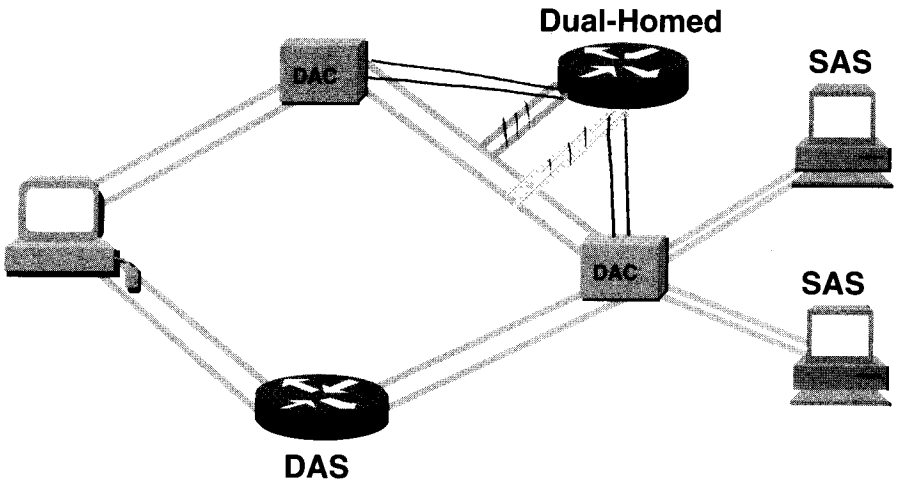

- **Devices attached to FDDI use token passing**

29

FDDI standards describe the physical layer and MAC sublayer.

Because FDDI specifies communication over fiber-optic cable, it is well suited for operations where nodes are separated by large distances or where networks must operate in electronically hostile environments such as factory floors.

FDDI has high speeds that make it suitable for network applications requiring large bandwidth—for example, video and graphics applications.

FDDI uses a token-passing protocol that operates on dual counter-rotating rings, as shown in the graphic. Under normal operation, data flows on a primary ring, while the secondary ring is idle. Some stations known as dual attachment stations (DASs) attach to both rings. Single attachment stations (SASs) have only a single physical medium dependent (PMD) connection to the primary ring by way of a dual-attached concentrator (DAC).

Mission-critical stations such as routers or mainframe hosts can use a technique called dual homing to provide additional fault-tolerance and help guarantee operation. With dual homing, a station is single-attached to two DACs, thereby providing an active primary link and a backup path to the FDDI LAN.

Copyright © 1996, Cisco Systems, Inc.

The FDDI Interface

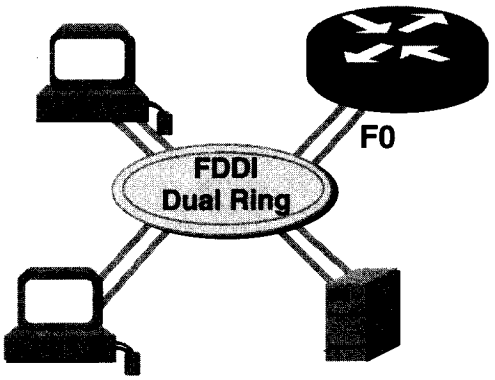

- **Cisco router's data link to FDDI uses an interface named *F* plus a number (for example, *F0*)**

30

FDDI is logically and physically a ring topology. Although it operates at higher speeds, FDDI is similar to Token Ring. The two network types share many features such as token passing and predictable deterministic delays.

All FDDI LAN stations use MAC addresses, including the router shown on the right of the graphic. The FDDI frame format uses four-bit symbols rather than eight-bit octets. Thus, the 48-bit MAC address for FDDI has 12 four-bit symbols.

For indicating the FDDI interface on the router, you will use the Cisco IOS interface type abbreviation *F* followed by an interface number (for example, *0,* as shown in the graphic).

Copyright © 1996, Cisco Systems, Inc.

FDDI Dual-Ring Reliability

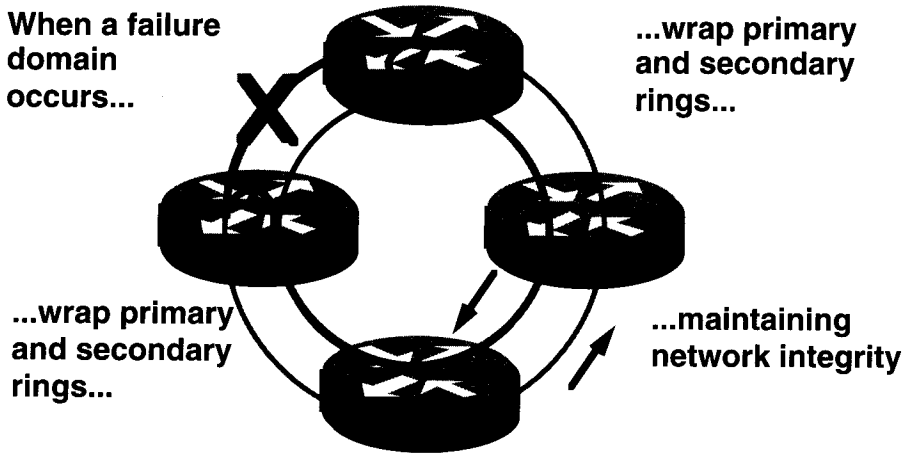

31

Access to the FDDI dual ring is determined by token possession. However, stations attach new tokens to the ends of their transmissions, and a downstream station is allowed to add its frame to the existing frame. Thus, at any given time, several information frames can be circling the ring.

All stations monitor the ring for invalid conditions such as a lost token, persistent data frames, or a break in the ring. If a node determines that no tokens have been received from its nearest active upstream neighbor (NAUN) during a predetermined time period, it begins transmitting beacon frames to identify the failure and its domain.

If a station receives its own beacon from upstream, it assumes that the ring has been repaired. If beaconing continues beyond a certain time limit, DASs on both sides of the failure domain loop (or wrap) the primary ring to the secondary ring to maintain network integrity (as illustrated in the graphic).

Copyright © 1996, Cisco Systems, Inc.

Exercise: Common LAN Technologies

Objective: Describe the topology and functionality of LANs.

Complete the following table to contrast Ethernet, Token Ring, and FDDI LAN technologies.

	Ethernet	Token Ring	FDDI
Cable and speed			
Router interface abbreviation (before the number)			

Copyright © 1996, Cisco Systems, Inc.

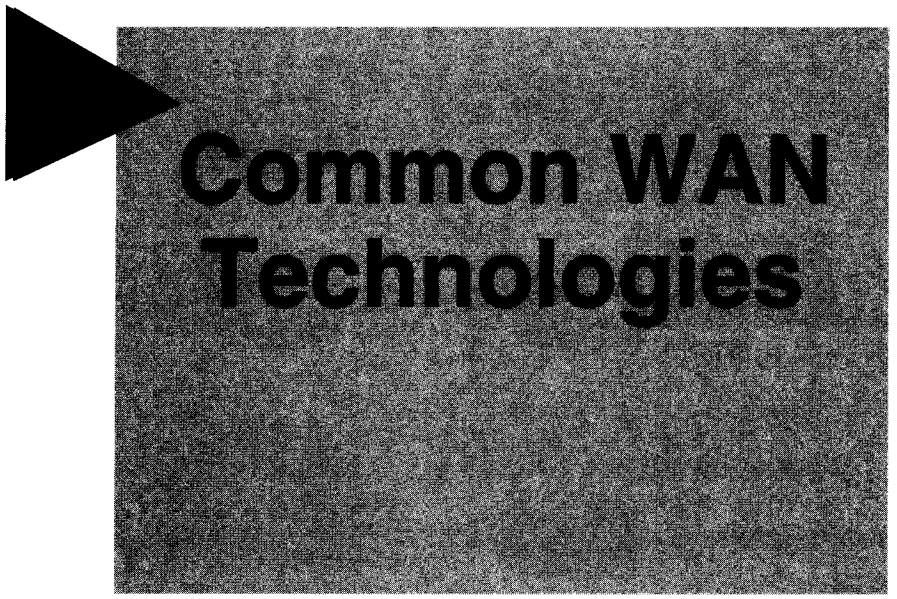

33

Common WAN Technologies

Copyright © 1996, Cisco Systems, Inc.

WAN Technology Overview

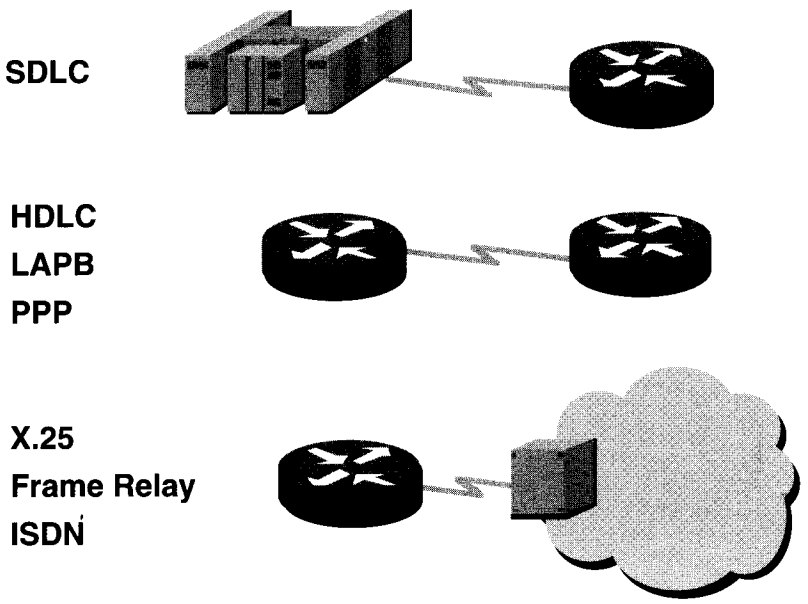

34

WAN physical layer protocols describe how to provide electrical, mechanical, operational, and functional connections for wide-area networking services. These services are most often obtained from WAN service providers such as Regional Bell Operating Companies (RBOCs), alternate carriers, and Post, Telephone, and Telegraph (PTT) agencies.

WAN data-link protocols describe how frames are carried between systems on a single data link. They include protocols designed to operate over dedicated point-to-point facilities, multipoint facilities based on dedicated facilities, and multiaccess switched services such as Frame Relay.

WAN standards are defined and managed by a number of recognized authorities including the following agencies:

- International Telecommunication Union-Telecommunication Standardization Sector (ITU-T), formerly the Consultative Committee for International Telegraph and Telephone (CCITT)
- International Organization for Standardization (ISO)
- Internet Engineering Task Force (IETF)
- Electronic Industries Association (EIA)

WAN standards typically describe both physical layer and data link layer requirements.

The graphic identifies several popular WAN services used in internetworks today. For example, ISDN integrates voice and data services on digital facilities. ISDN has grown as a preferred facility for accessing World Wide Web (WWW) multimedia.

Copyright © 1996, Cisco Systems, Inc.

Physical Layer: WAN

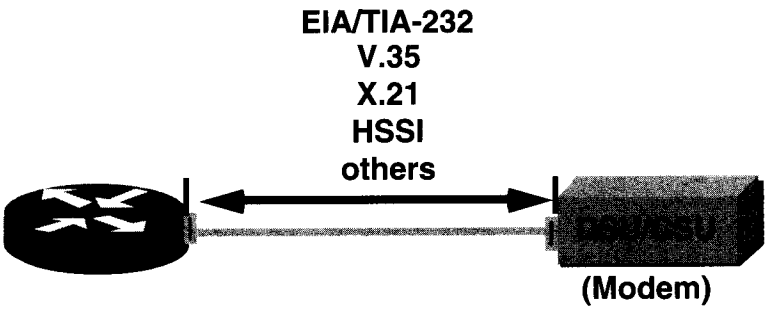

DTE
Data Terminal Equipment
End of the user's device on the WAN link

DCE
Data Circuit-Terminating Equipment
End of the WAN provider's side of the communication facility

35

The WAN physical layer describes the interface between the data terminal equipment (DTE) and the data circuit-terminating equipment (DCE). Typically, the DCE is the service provider, and the DTE is the attached device. In this model, the services offered to the DTE are made available through a modem or channel service unit/data service unit (CSU/DSU).

Several physical layer standards specify this interface:

- EIA/TIA-232
- EIA/TIA-449
- V.24
- V.35
- X.21
- G.703
- EIA-530
- High-Speed Serial Interface (HSSI)

Copyright © 1996, Cisco Systems, Inc.

Data Link Layer: WAN Protocols

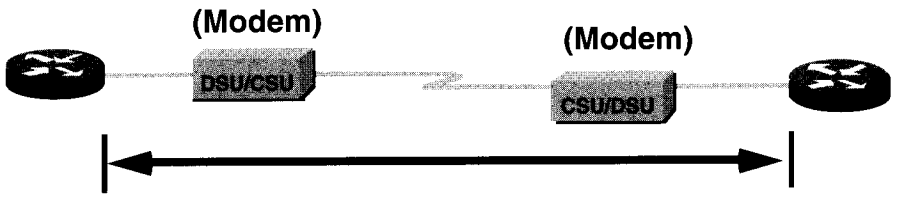

- **SDLC—Synchronous Data Link Control**
- **HDLC—High-Level Data Link Control**
- **LAPB—Link Access Procedure Balanced**
- **Frame Relay—Simplified version of HDLC framing**
- **PPP—Point-to-Point Protocol**
- **X.25—Packet level protocol (PLP)**
- **ISDN—Integrated Services Digital Network (data-link signaling)**

36

The common data-link encapsulations associated with synchronous serial lines are listed in the graphic:

- Synchronous Data Link Control (SDLC)—A bit-oriented protocol developed by IBM. SDLC defines a multipoint WAN environment that allows several stations to connect to a dedicated facility. SDLC defines a primary station and one or more secondary stations. Communication is always between the primary station and one of its secondary stations. Secondary stations cannot communicate with each other directly.
- High-Level Data Link Control (HDLC)—An ISO standard. HDLC might not be compatible between different vendors because of the way each vendor has chosen to implement it. HDLC supports both point-to-point and multipoint configurations.
- Link Access Procedure, Balanced (LAPB)—Primarily used with X.25, but can also be used as a simple data-link transport. LAPB includes capabilities for detecting out-of-sequence or missing frames as well as for exchanging, retransmitting, and acknowledging frames.
- Frame Relay—Uses high-quality digital facilities where the error checking of LAPB is unnecessary. By using a simplified framing with no error correction mechanisms, Frame Relay can send Layer 2 information very rapidly, compared to these other WAN protocols.
- Point-to-Point Protocol (PPP)—Described by RFC 1661, two standards developed by the IETF. PPP contains a protocol field to identify the network-layer protocol.
- X.25—Defines the connection between a terminal and a packet-switching network.
- Integrated Services Digital Network—A set of digital services that transmits voice and data over existing phone lines.

Copyright © 1996, Cisco Systems, Inc.

Summary

The physical layer provides access to the wires of an internetwork

The data link layer provides support for communication over several types of data links:

- **LAN (Ethernet/IEEE 802.3, Token Ring/IEEE 802.5, FDDI)**
- **Dedicated WAN (SDLC, HDLC, PPP, LAPB)**
- **Switched WAN (X.25, Frame Relay, ISDN)**

37

Copyright © 1996, Cisco Systems, Inc.

Exercise: Physical and Data Link Layers

Problem 1

Objective: Identify and describe the data link sublayers and their functions.

In the diagram, fill in the layer and sublayer names of the OSI model on the space provided inside each rectangle. Then on the lines provided, write a brief description of the functions performed by sublayers A and B.

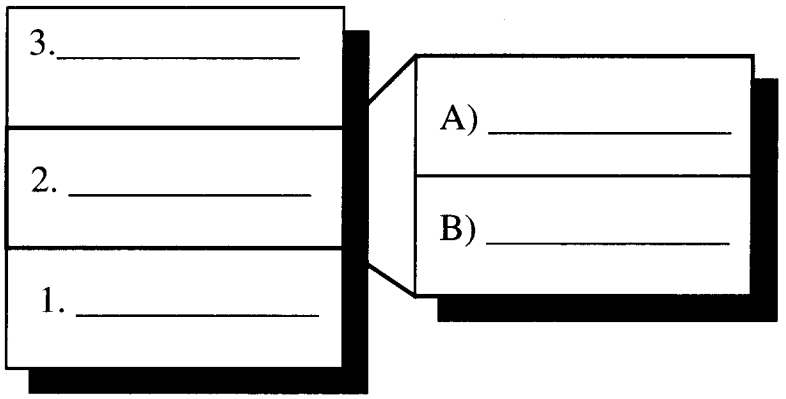

A)

B)

Problem 2

Refer to the diagram in problem 1. Assume you have a logical address used by the function in Layer 3. Name a method used to find a device address used by the function in Layer 2. Write your answer on the line below.

Copyright © 1996, Cisco Systems, Inc.

Problem 3

Objective: Differentiate between LAN and WAN protocols.

Objective: Describe the characteristics of WAN-based protocols.

Entries in the table below list the names of a LAN or WAN protocol or standard in column 1. Column 2 statements summarize a LAN or WAN topology, function, or characteristic. In the left column (labeled Write Letters) write two answer letters in each space:

- Write *L* if the protocol or standard is used on LANs, or write a *W* if it is used on WANs.
- Also write the letter identifying the correct statement in column 2 that describes the given protocol or standard.

For example, for item 2, write *W* followed by *C* in the blank space if you think item 2's protocol or standard is used on a WAN and summarized in item *C*.

Note Column 2 intentionally contains extra statements.

Write Letters	Column 1 Protocol or Standard	Column 2 Topology, Function, or Characteristic
1. _____	SDLC	A) Equivalent of IEEE 802.5
2. _____	EIA/TIA-232	B) Either side can initiate connection to access link
3. _____	802.3	C) Voice-grade access, formerly a recommended standard
4. _____	Frame Relay	D) Uses primary and secondary roles for IBM data links
5. _____	Ethernet II	E) From original Xerox work; uses field for protocol type
6. _____	ISDN	F) Emerging standard to provide TCP/IP internet access
7. _____	HDLC	G) Proprietary version of a default protocol for Cisco routers
8. _____	Token Ring	H) From IEEE efforts; uses field for length rather than type
9. _____	FDDI	I) Uses simplified HDLC for higher-speed communication
		J) Uses four-bit symbols rather than octets in its framing
		K) Integrates voice and data services on digital facilities

Copyright © 1996, Cisco Systems, Inc.

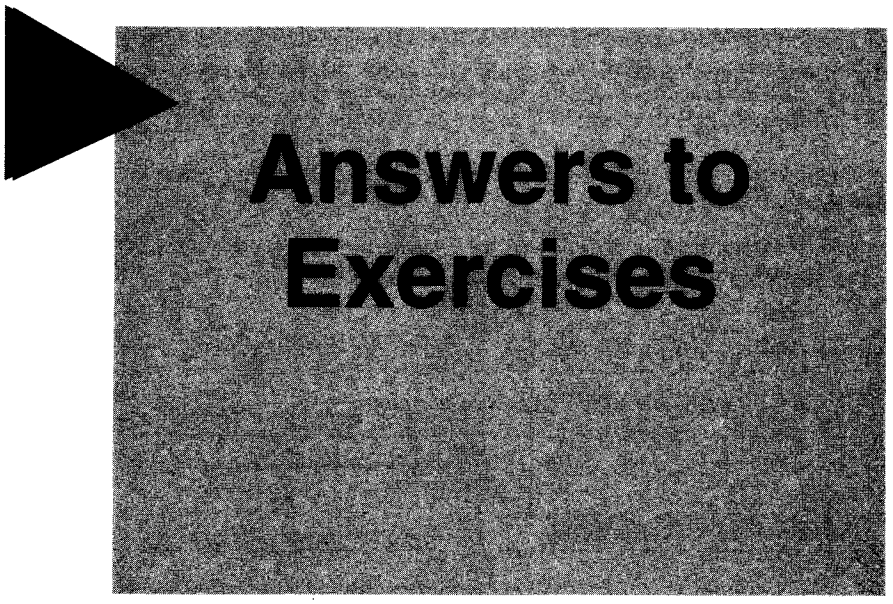

40

Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

Exercise: MAC Addresses

Host A addresses:

for host B 0000.0c12.2222

for host D 0000.0c12.3333

Host B addresses:

for host A 0000.0c12.1111

for host D 0000.0c12.3333

Host D addresses:

for host A 0000.0c12.3333

for host B 0000.0c12.3333

Router C addresses:

for host A 0000.0c12.1111

for host B 0000.0c12.2222

for host D 0000.0c12.4444

Exercise: Common LAN Technologies

	Ethernet	Token Ring	FDDI
Cable and speed	10Base2, 10Base5, 10BaseT, 100BaseTX, 100BaseT4, 100BaseFX, and 100BaseVG	Not specified in 802.5; normally 4 or 16 Mbps	Fiber optic; 100 Mbps
Router interface abbreviation (before the number)	E	To	F

Copyright © 1996, Cisco Systems, Inc.

Exercise: Physical and Data Link Layers

Problem 1

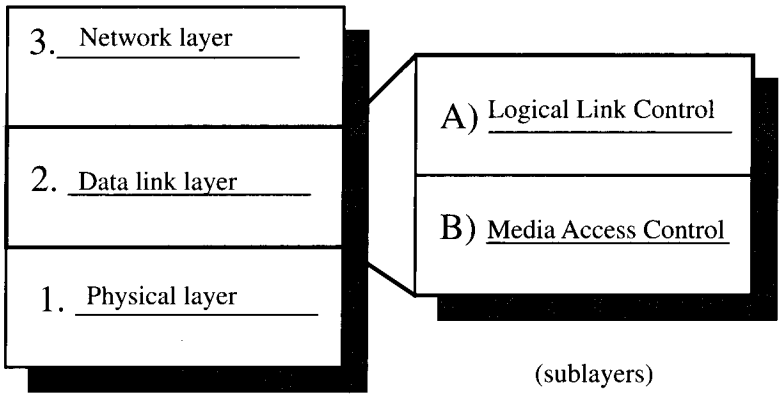

Here are some of many possible correct answers for the fill-in-the-blank-lines part of this problem.

A) LLC sublayer—Manages communication between devices over a single data link using fields allowing higher-layer software to share the use of the data link. It uses SAPs as interfaces that MAC sublayer functions can use. It offers optional connection, flow control, and sequencing services.

B) MAC sublayer—Provides source and destination addresses for data-link frames, provides a 48-bit address for identifying LAN devices (or their ports), and gives vendors an approved numbering scheme for serializing their devices. This sublayer can also offer a unique identifier for a device that can be modified (for example, by DECnet) or software that can be configured (for example, by VTAM).

Problem 2

ARP—Address Resolution Protocol.

Note: Not all protocol suites require this Layer 3-to-Layer-2 function.

Problem 3

1. W D
2. W C
3. L H
4. W I
5. L E
6. W K
7. W G
8. L A
9. L J

Copyright © 1996, Cisco Systems, Inc.

CHAPTER 4

Network Layer and Path Determination

Copyright © 1996, Cisco Systems, Inc.

Objectives

Upon completion of this chapter, you will be able to:

List the key internetworking functions of the OSI network layer and how they are performed in a router

Describe the two parts of network addressing, then identify the parts in specific protocol address examples

Contrast the network discovery and update processes in distance vector routing with those in link-state routing

List problems that each routing type encounters when dealing with topology changes, and describe techniques to reduce the number of these problems

Explain the services of separate and integrated multiprotocol routing

2

This chapter discusses the network layer of the OSI reference model. It covers basic information such as how network-layer addressing works with different protocols. It explains the difference between routing and routed protocols and contrasts static and dynamic routes. It explains how routers track the distance between locations.

The chapter then covers distance vector, link-state, and hybrid routing approaches. It explains the strengths of each approach and describes how each resolves common routing problems.

Sections:

- Network Layer Basics
- Routing Protocols
- Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

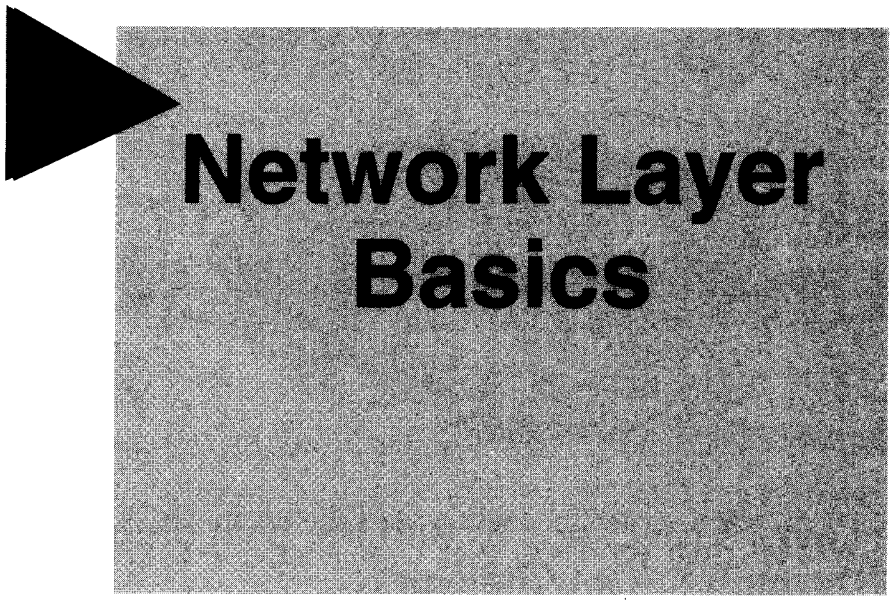

3

Network Layer Basics

Copyright © 1996, Cisco Systems, Inc.

Network Layer: Path Determination

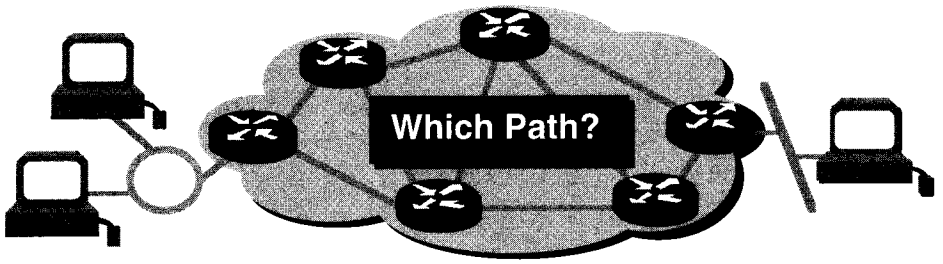

- **Layer 3 functions to find the best path through the internetwork**

4

Which path should traffic take through the cloud of networks? Path determination occurs at Layer 3, the network layer. The path determination function enables a router to evaluate the available paths to a destination and to establish the preferred handling of a packet.

Routing services use network topology information when evaluating network paths. This information can be configured by the network administrator or collected through dynamic processes running in the network.

The network layer interfaces to networks and provides best effort end-to-end packet delivery services to its user, the transport layer. The network layer sends packets from the source network to the destination network.

After the router determines which path to use, it can proceed with switching the packet: taking the packet it accepted on one interface and forwarding it to another interface or port that reflects the best path to the packet's destination.

Copyright © 1996, Cisco Systems, Inc.

Network Layer: Communicate Path

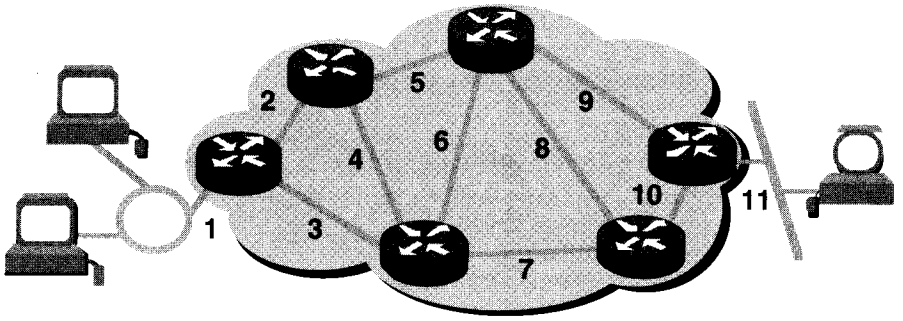

- **Addresses represent the path of media connections**

5

To be truly practical, an internetwork must consistently represent the paths of its media connections. As the graphic shows, each line between the routers has a number that the routers use as a network address. These addresses must convey information that can be used by a routing process. This means that an address must have information about the path of media connections used by the routing process to pass packets from a source toward a destination.

The network layer combines this information about the path of media connections—sets of links—into an internetwork by adding path determination, path switching, and route processing functions to a communication system. Using these addresses, the network layer also provides a relay capability that interconnects independent networks.

The consistency of Layer 3 addresses across the entire internetwork also improves the use of bandwidth by preventing unnecessary broadcasts. Broadcasts invoke unnecessary process overhead and waste capacity on any devices or links that do not need to receive the broadcast.

By using consistent end-to-end addressing to represent the path of media connections, the network layer can find a path to the destination without unnecessarily burdening the devices or links on the internetwork with broadcasts.

Copyright © 1996, Cisco Systems, Inc.

Addressing: Network and Host

Network	Host
1	1 2 3
2	1
3	1

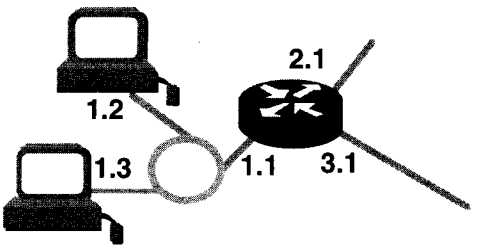

- **Network address—Path part used by the router**
- **Host address—Specific port or device on the network**

6

The network address identifies a path part used by the router within the internetwork cloud. The router uses the network address to identify the source or destination network of a packet within an internetwork. The graphic shows three network numbers emanating from the router.

For some network-layer protocols, this relationship is established by a network administrator who assigns network addresses according to some preconceived internetwork addressing plan. For other network-layer protocols, assigning addresses is partially or completely dynamic.

Most network-protocol addressing schemes use some form of a host or node address. The host address refers to the device's specific port or device on the network. For instance, in the graphic three hosts are shown sharing the network number 1. The host or node address identifies that the packet is on its source or destination port or device on the network. For LANs, this port or device address can reflect the real Media Access Control (MAC) address of the device.

However, unlike a MAC address that has a preestablished and usually fixed relationship to a device, a network address has a logical relationship.

Copyright © 1996, Cisco Systems, Inc.

Protocol Addressing Variations

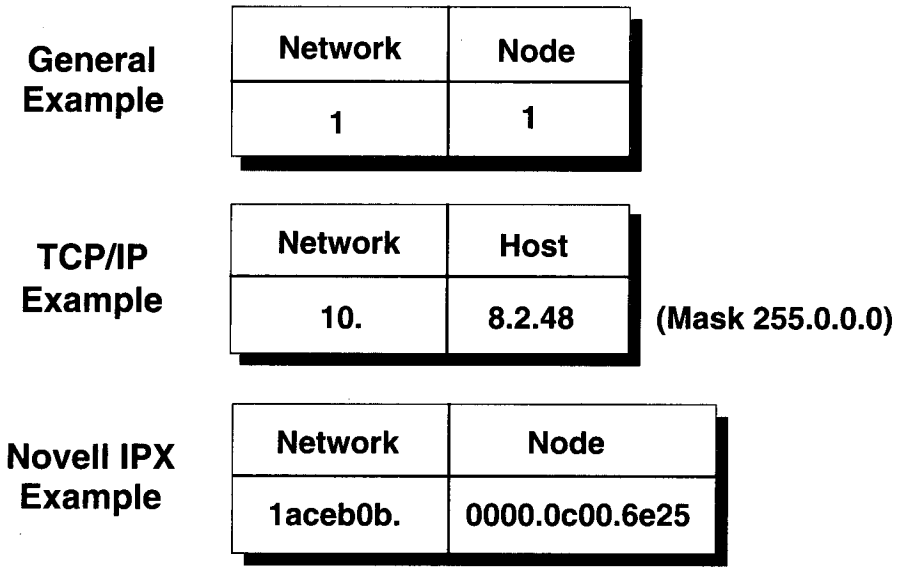

7

The two-part network addressing scheme extends across all the protocols covered in this course. How do you interpret the meaning of the address parts? What authority allocates the addresses? These answers vary from protocol to protocol.

For example, in the TCP/IP example IP address, dotted decimal numbers show a network part and a host part. The network 10 uses the first of the four numbers as the network part and the last three sets of numbers—8.2.48—as a host address. The mask is a companion number to the IP address. It communicates to the router the part of the number to interpret as the network number and identifies the remainder available for host addresses inside that network.

The Novell IPX example uses a different variation of this two-part address. The network address 1aceb0b is a hexadecimal (base 16) number that cannot exceed a fixed maximum number of digits. The host address 0000.0c00.6e25 (also a hexadecimal number) is a fixed 48 bits long. This host address derives automatically from information in the hardware of the specific LAN device.

These are the two most common Layer 3 address types. You will learn more about these and other protocol addressing rules in the next few pages. Then you will use valid Layer 3 addresses during the hands-on labs later in this course.

Copyright © 1996, Cisco Systems, Inc.

TCP/IP Network Addressing

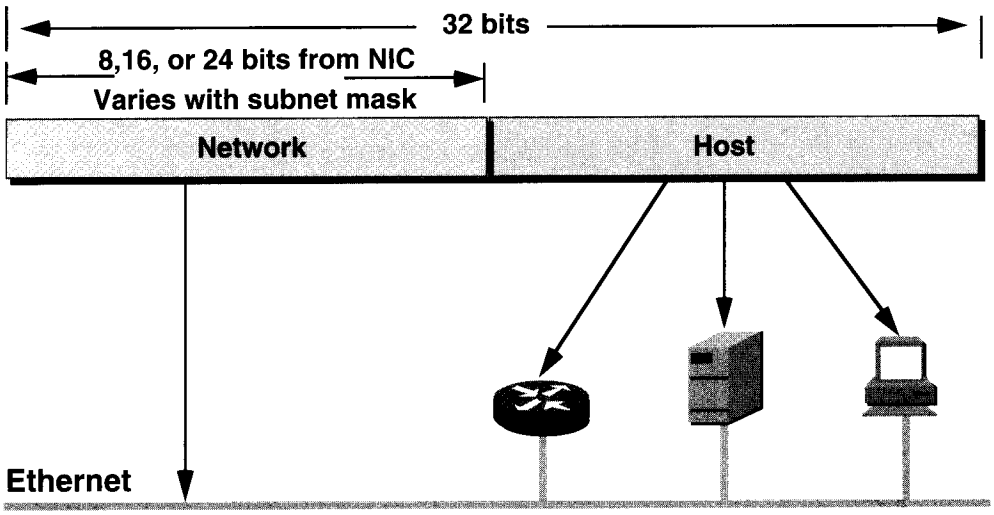

8

TCP/IP networks represent addresses as 32-bit entities, divided into a network portion and a host portion. The Internet Request For Comments (RFC) 1020 divides the network portion into classes. All classes of specific, Internet-legal network addresses come from a central authority: the Network Information Center (NIC). The most common of these classes follow:

- Class A—Using 8 bits for the network, with the remaining 24 bits for host addressing
- Class B—Using 16 bits for the network, with the remaining 16 bits for host addressing
- Class C—Using 24 bits for the network, with the remaining 8 bits for host addressing
- Class D—Used for IP multicast addresses

IP networks typically are subdivided into subnetworks. When an IP address has been subnetted, the network part of the address is described by two elements: the network number, still assigned by the NIC, and the subnetwork number, assigned by the local network administrator.

Copyright © 1996, Cisco Systems, Inc.

Other Protocol Addressing

Protocol	Network Address	Node Address
Novell IPX	Up to 32 bits (hex); refers to the media (for example, Ethernet)	48 bits (hex); usually MAC address of a LAN interface
AppleTalk	Up to 16 bits (dec); refers to one, or one of many nets in cable range on media	Up to 8 bits added to network number; usually dynamically assigned on LAN
X.25 (X.121)	4 (dec) digits of DNIC with 2- or 3-digit Data Country Code and 1 network digit	Up to 10 or 11 digits of Network Terminal Number; usually from WAN service provider

- **Several other Layer 3 addressing schemes**

9

A router can handle many other protocol addressing schemes. The table and text summarize key details about three of the most common of these:

- Novell Internet Packet Exchange (IPX) address—Identifies the IPX network in the first part of the address using an address range of 32 bits to reflect a hexadecimal number. This network number refers to the medium, for example, the Ethernet or Token Ring LAN. For the node address, IPX uses a 48-bit hexadecimal number, usually derived automatically from the MAC address of a LAN interface to the IPX network.
- AppleTalk address—Identifies the network in the first part of the address. The 16-bit network numbers are assigned to physical links either individually or in ranges called cable ranges. This approach makes it possible for many network addresses to use the same LAN media. The 8-bit AppleTalk node portion is called the host address. An Apple end station usually acquires this host address dynamically when it boots up onto the network.
- X.25 address—Within the X.25 protocol suite, the X.121 protocol covers the international numbering plan for public data networks (PDNs). The network portion of the address specifies three or four decimal digits as the Data Network Identification Code (DNIC). This DNIC includes a Data Country Code (DCC). An example is 310 for the United States, followed by the network number 6 for Tymnet—one of the major PDNs. The node address portion is called the network terminal number (NTN). X.25 users usually obtain these NTNs from an authority within the X.25 data network service provider.

Cisco routers can handle these and many other protocol-specific Layer 3 addressing schemes.

Copyright © 1996, Cisco Systems, Inc.

Routing Uses Network Addresses

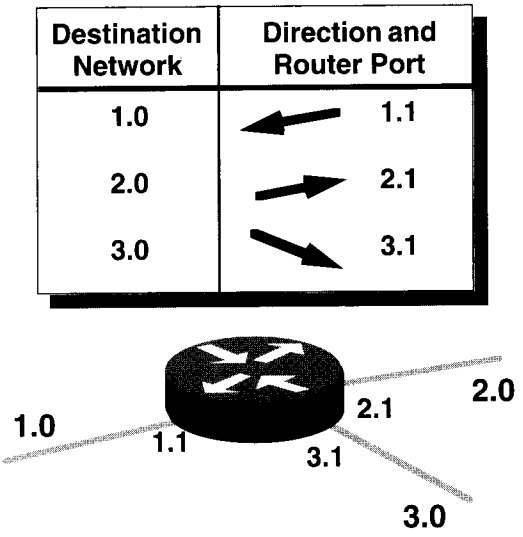

- **Network portion of address used to make path selections**
- **Node portion of address refers to router port to the path**

10

Routers generally relay a packet from one data link to another. To relay a packet, a router uses two basic functions: a path determination function and a switching function.

The graphic illustrates how routers use the addressing for routing and switching functions.

Although the path determination function sometimes is able to calculate the complete path from the router to the destination, a router is responsible only for passing the packet to the best network along the path. This best path is represented as a direction to a destination network—like the arrows in the figure pointing to the next hop. The router uses the network portion of the address to make path selections.

The switching function allows a router to accept a packet on one interface and forward it on a second interface. The path determination function enables the router to select the most appropriate interface for forwarding a packet. The node portion of the address refers to a specific port on the router that leads to an adjacent router in that direction.

Copyright © 1996, Cisco Systems, Inc.

Routed versus Routing Protocol

- **Routed protocol used between routers to direct user traffic**

  Examples: IP, IPX

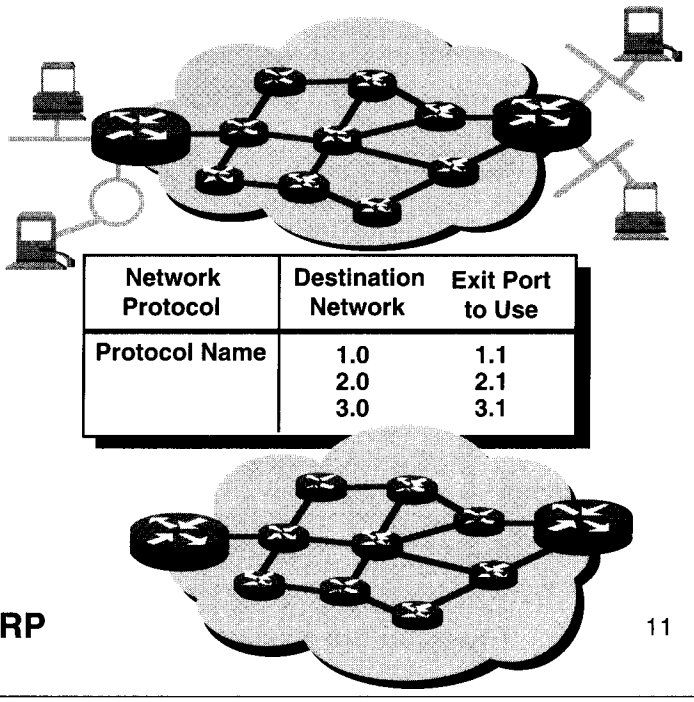

- **Routing protocol used only between routers to maintain tables**

  Examples: RIP, IGRP

11

Confusion often exists between the similar terms rout*ing* protocol and rout*ed* protocol:

- Routed protocols—Any network protocol suite that provides enough information in its network layer address to allow a packet to direct user traffic. Routed protocols define the format and use of the fields within a packet. Packets generally are conveyed from end system to end system. The Internet protocol IP and Novell's IPX are examples of routed protocols.
- Routing protocol—Supports a rout*ed* protocol by providing mechanisms for sharing routing information. Routing protocol messages move between the routers. A routing protocol allows the routers to communicate with other routers to update and maintain tables. Routing protocol messages do not carry end-user traffic from network to network. A routing protocol uses the routed protocol to pass information between routers. TCP/IP examples of routing protocols are Routing Information Protocol (RIP), Interior Gateway Routing Protocol (IGRP), and Open Shortest Path First (OSPF).

Copyright © 1996, Cisco Systems, Inc.

Network-Layer Protocol Operations

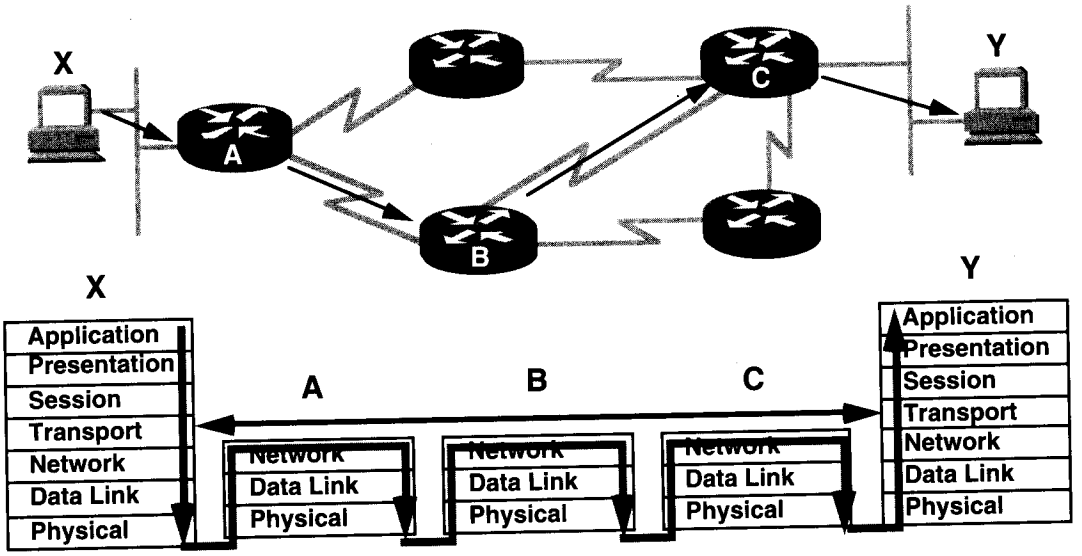

- **Each router provides its services to support upper-layer functions**

12

When a host application needs to send a packet to a destination of a different network, a data-link frame is received on one of a router's interfaces. The router decapsulates and examines the frame to determine what type of network-layer data is being carried. The network-layer data is sent to the appropriate network-layer process, and the frame itself is discarded.

The network layer process examines the header to determine the destination network and then references the routing table that associates networks to outgoing interfaces.

The packet is again encapsulated in the data-link frame for the selected interface and queued for delivery to the next hop in the path.

This process occurs each time the packet switches through another router. At the router connected to the network containing the destination host, the packet is again encapsulated in the destination LAN's data-link frame type for delivery to the protocol stack on the destination host.

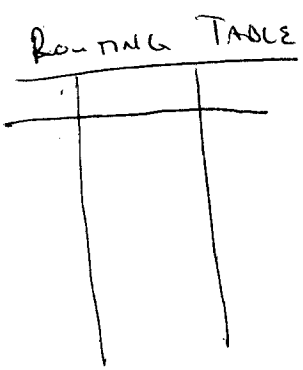

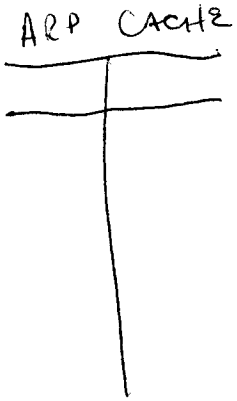

Copyright © 1996, Cisco Systems, Inc.

Multiprotocol Routing

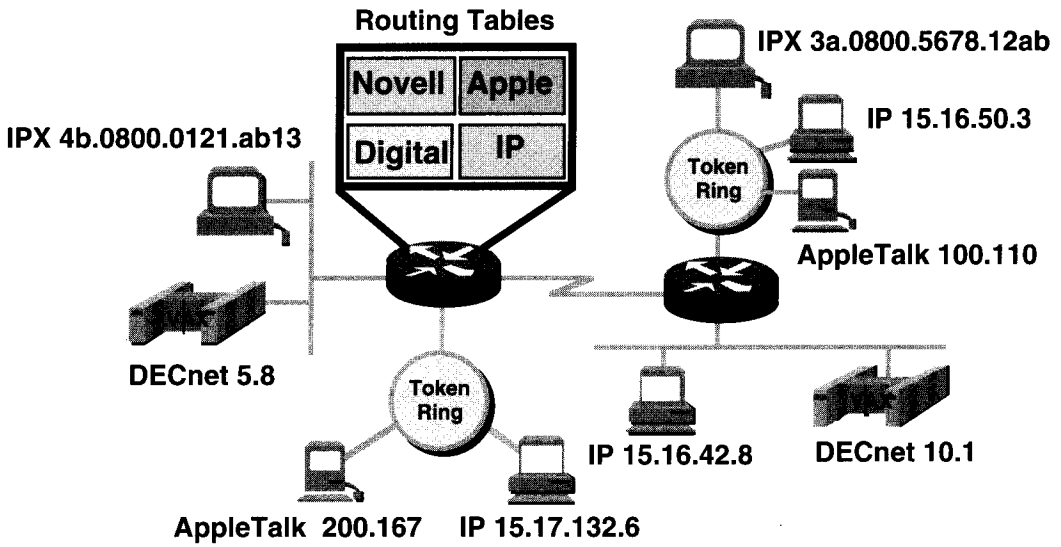

- **Routers pass traffic from all routed protocols over the internetwork**

13

Routers are capable of supporting multiple independent routing algorithms and maintaining associated routing tables for several routed protocols concurrently. This capability allows a router to interleave packets from several routed protocols over the same data links.

The routed and routing protocols have no knowledge of other protocols. This concept is called ships-in-the-night routing.

Copyright © 1996, Cisco Systems, Inc.

Static versus Dynamic Routes

Static Route

Uses a protocol route that a network administrator enters into the router

Dynamic Route

Uses a route that a network routing protocol adjusts automatically for topology or traffic changes

14

Static knowledge is administered manually: A network administrator enters it into the router's configuration. The administrator must manually update this static route entry whenever an internetwork topology change requires an update. Static knowledge can be private—by default it is not conveyed to other routers as part of an update process. You can, however, configure the router to share this knowledge.

Dynamic knowledge works differently. After the network administrator enters configuration commands to start dynamic routing, route knowledge is updated automatically by a routing process whenever new topology information is received from the internetwork. Changes in dynamic knowledge are exchanged between routers as part of the update process.

Copyright © 1996, Cisco Systems, Inc.

Static Route Example

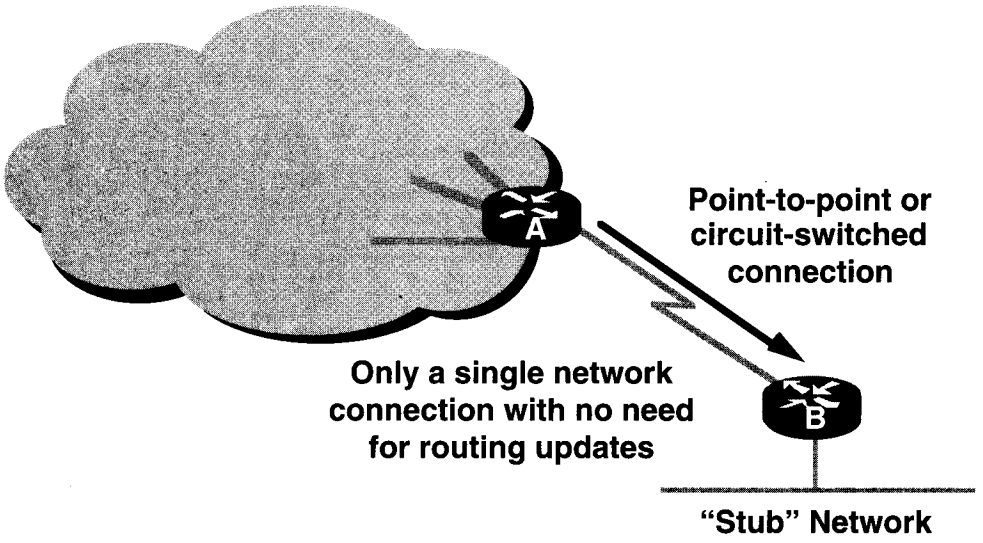

- **Fixed route to address reflects administrator's knowledge**

15

When it reflects a network administrator's special knowledge about network topology, static routing has several useful applications.

Dynamic routing tends to reveal everything known about an internetwork. For security reasons, it might be appropriate to conceal parts of an internetwork. Static routing allows an internetwork administrator to specify what is advertised about restricted partitions.

When an internetwork partition is accessible by only one path, a static route to the partition can be sufficient. This type of partition is called a stub network. Configuring static routing to a stub network avoids the overhead of dynamic routing.

Copyright © 1996, Cisco Systems, Inc.

Default Route Example

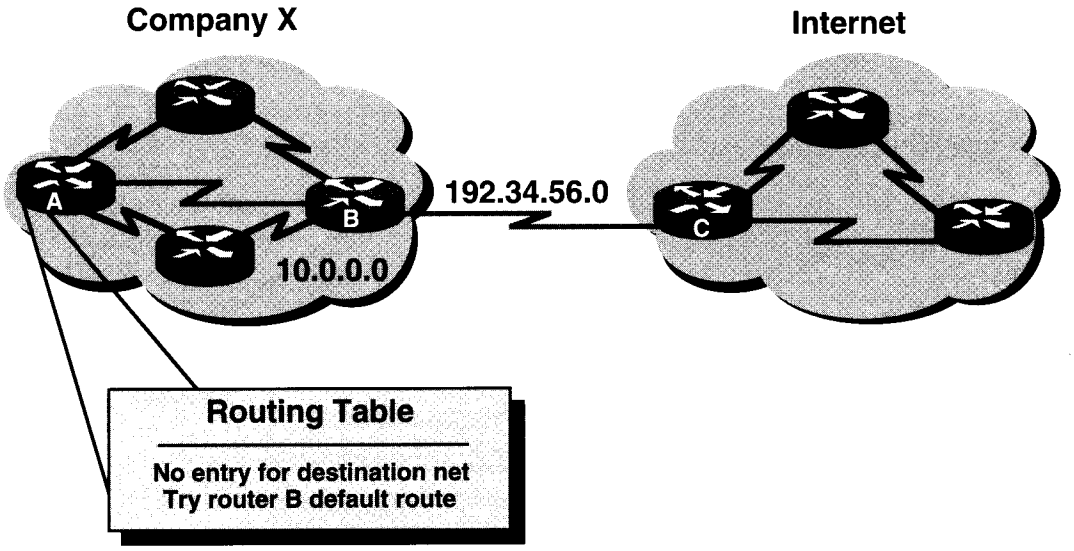

- **Use if next hop is not explicitly listed in the routing table** 16

The graphic shows a use for a default route—a routing table entry that is used to direct frames for which the next hop is not explicitly listed in the routing table. Default routes can be set as the result of the administrator's static configuration, or they can be set by dynamic routing with knowledge of many protocols.

In this example, Company X routers possess specific knowledge of the topology of the Company X internetwork, but not of other networks. Maintaining knowledge of every other internetwork accessible by way of the Internet cloud is unnecessary and unreasonable, if not impossible.

Instead of maintaining specific internetwork knowledge, each router in Company X is informed by the default route that it can reach any unknown destination by directing the packet to the Internet.

Copyright © 1996, Cisco Systems, Inc.

Adapting to Topology Change

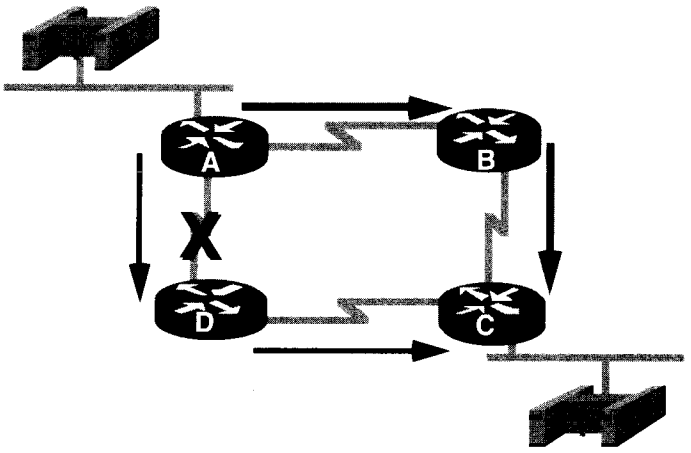

- **Can an alternate route substitute for a failed route?**

17

The internetwork shown in the graphic adapts differently to topology changes depending on whether it uses statically or dynamically configured knowledge.

Static knowledge allows the routers to properly route a packet from network to network. The router refers to its routing table and follows the static knowledge there to relay the packet to router D. Router D does the same and relays the packet to router C. Router C delivers the packet to the destination host.

But what happens if the path between router A and router D fails? Obviously router A will not be able to relay the packet to router D with a static route. Until router A is manually reconfigured to relay packets by way of router B, communication with the destination network is impossible.

Dynamic knowledge offers more automatic flexibility. According to the routing table generated by router A, a packet can reach its destination over the preferred route through router D. However, a second path to the destination is available by way of router B. When router A recognizes the link to router D is down, it adjusts its routing table, making the path through router B the preferred path to the destination. The routers continue sending packets over this link.

When the path between routers A and D is restored to service, router A can once again change its routing table to indicate a preference for the counterclockwise path through routers D and C to the destination network.

Copyright © 1996, Cisco Systems, Inc.

Dynamic Routing Operations

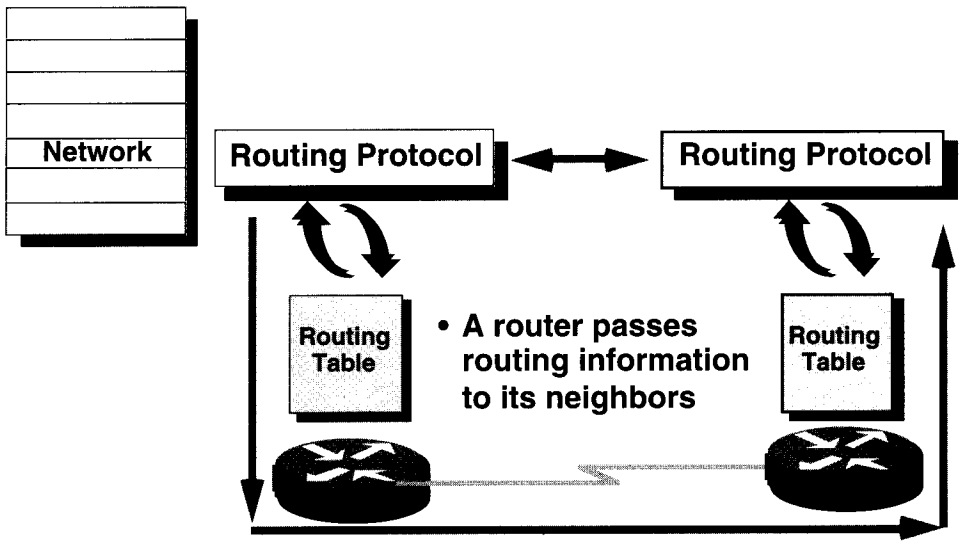

- **Routing protocol maintains and distributes routing information**

18

The success of dynamic routing depends on two basic router functions:

- Maintenance of a routing table
- Timely distribution of knowledge—in the form of routing updates—to other routers

Dynamic routing relies on a routing protocol to disseminate knowledge. A routing protocol defines the set of rules used by a router when it communicates with neighboring routers. For example, a routing protocol describes:

- How updates are conveyed
- What knowledge is conveyed
- When to convey knowledge
- How to locate recipients of the updates

Copyright © 1996, Cisco Systems, Inc.

Representing Distance with Metrics

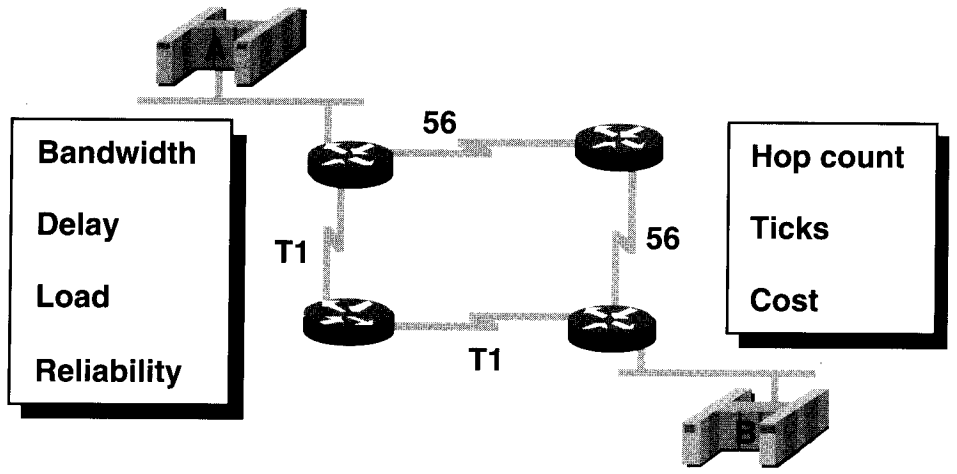

- **Information used to select the best path for routing**

19

When a routing algorithm updates the routing table, its primary objective is to determine the best information to include in the table. Each routing algorithm interprets best in its own way. The algorithm generates a number—called the metric value—for each path through the network. Typically, the smaller the metric, the better the path.

Metrics can be calculated based on a single characteristic of a path. You can calculate more complex metrics by combining several characteristics. Several path characteristics are used in metric calculations. The metrics most commonly used by routers follow:

- Bandwidth—Data capacity of a link. For instance, normally, a 10-Mbps Ethernet link is preferable to a 64-kbps leased line.
- Delay—Length of time required to move a packet from source to destination.
- Load—Amount of activity on a network resource such as a router or link.
- Reliability—Usually refers to the bit-error rate of each network link.
- Hop count—Number of passages of a packet through the output port of one router.
- Ticks—Delay on a data link using IBM PC clock ticks (approximately 55 milliseconds).
- Cost—Arbitrary value, usually based on bandwidth, dollar expense, or other measurement, that is assigned by a network administrator.

Copyright © 1996, Cisco Systems, Inc.

Exercise: Network Layer Basics

Problem 1

Objective: List the key internetworking functions of the OSI network layer and how they are performed in a router.

Identify three key network-layer functions and briefly describe how a router performs these functions.

Problem 2

Objective: Describe the two parts of network addressing, then identify the parts in specific protocol address examples.

Fill in the blanks to complete the following statement:

The two general parts of a Layer 3 address are a _____ part and a _____ part.

Problem 3

Identify the protocol suite and the two parts of the common Layer 3 addresses shown:

A) 131.108.3.1 (assume a subnet mask of 255.255.0.0)

B) 1000.128 (assume a cable range of 1000-1000)

C) abadcafe.0000.0c56.de33

D) 31060004085551

Copyright © 1996, Cisco Systems, Inc.

Problem 4

Fill in the blanks:

A) A router keeps location information in a

_____.

B) The protocol that routers use to communicate locations to one another is a

_____.

C) How many tables are used in multiprotocol routing?

_____.

Copyright © 1996, Cisco Systems, Inc.

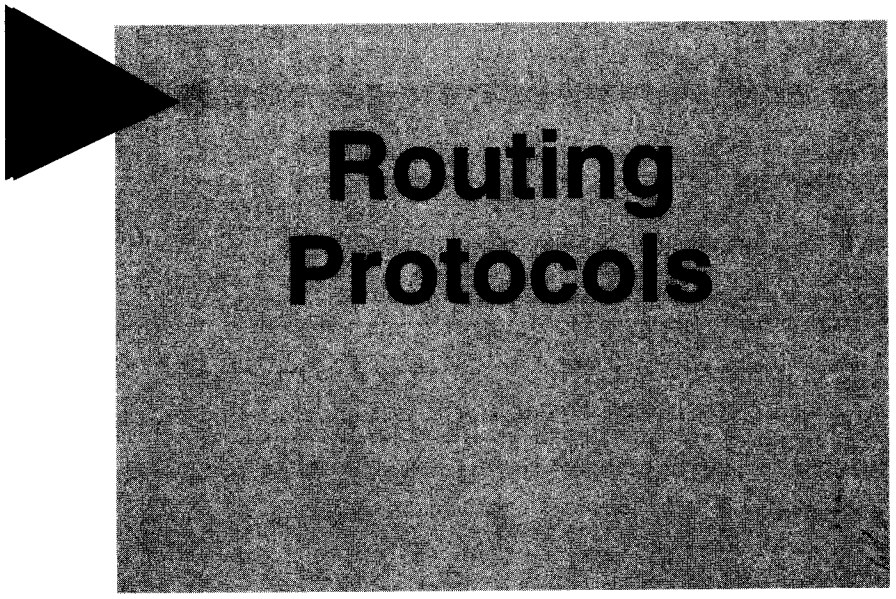

22

Routing Protocols

Copyright © 1996, Cisco Systems, Inc.

Classes of Routing Protocols

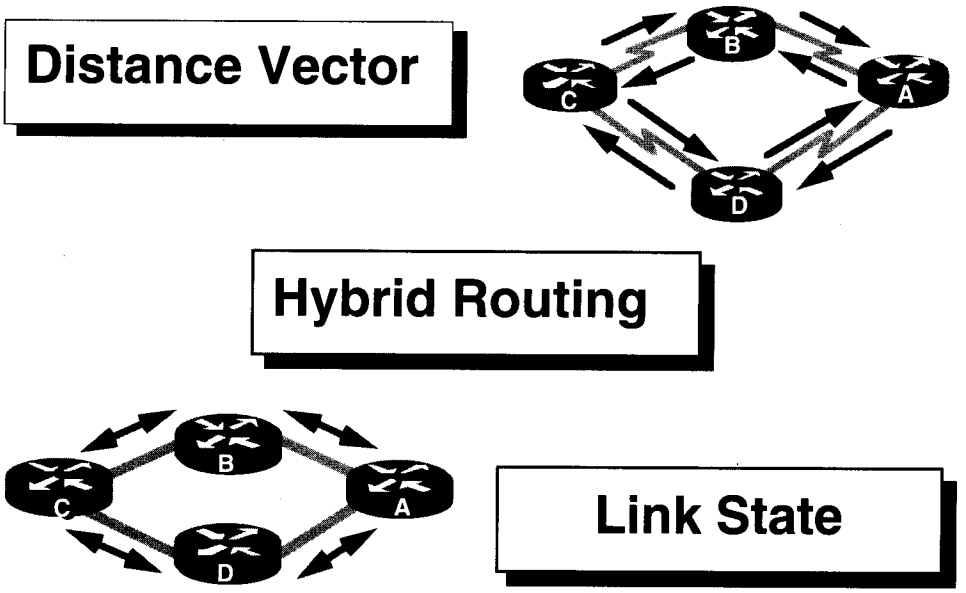

Most routing algorithms can be classified as conforming to one of two basic algorithms: distance vector or link state.

The distance vector routing approach determines the direction (vector) and distance to any link in the internetwork.

The link-state (also called shortest path first) approach re-creates the exact topology of the entire internetwork (or at least the partition in which the router is situated).

The balanced hybrid approach combines aspects of the link-state and distance vector algorithms.

The next several pages cover procedures and problems for each of these routing algorithms and present techniques for minimizing the problems.

There is no single best routing algorithm for all internetworks. Network administrators must weigh technical and nontechnical aspects of their network to determine the best algorithm. Cisco IOS software can configure whatever routing choices best fit the administrator's internetwork.

Copyright © 1996, Cisco Systems, Inc.

One Issue: Time to Convergence

Convergence occurs when all routers use a consistent perspective of network topology

After a topology changes, routers must recompute routes, which disrupts routing

The process and time required for router reconvergence varies in routing protocols

24

The routing algorithm is fundamental to dynamic routing. Whenever the topology of the internetwork changes because of growth, reconfiguration, or failure, the internetwork knowledge base must also change.

The knowledge base needs to reflect an accurate, consistent view of the new topology. This accurate, consistent view is called convergence. When all routers in an internetwork are operating with the same knowledge, the internetwork is said to have converged.

Fast convergence is a desirable internetwork feature because it reduces the period of time that routers have outdated knowledge for making routing decisions that could be incorrect, wasteful, or both.

Copyright © 1996, Cisco Systems, Inc.

Distance Vector Concept

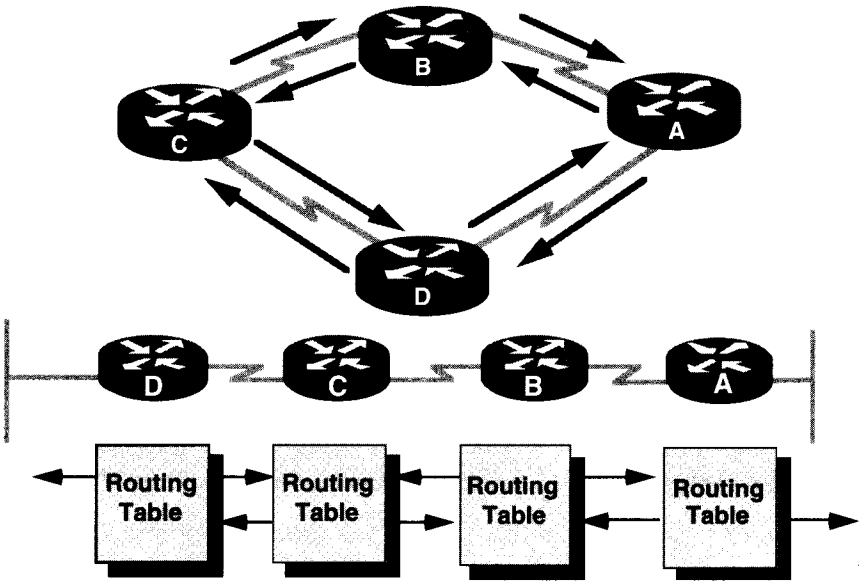

- **Pass periodic copies of routing table to neighbor routers and accumulate distance vectors**

25

Distance vector-based routing algorithms (also known as Bellman-Ford algorithms) pass periodic copies of a routing table from router to router. Regular updates between routers communicate topology changes.

Each router receives a routing table from its direct neighbor. For example, in the graphic, router B receives information from router A. Router B adds a distance vector number (such as a number of hops) increasing the distance vector, then passes the routing table to its other neighbor, router C. This same step-by-step process occurs in all directions between direct-neighbor routers.

In this way, the algorithm accumulates network distances so it can maintain a database of internetwork topology information. Distance vector algorithms do not allow a router to know the exact topology of an internetwork.

RIP - IP, IPX, XNS
IGRP - IP
RTMP - AT
RTP - Vines
Dec Phase IV

Copyright © 1996, Cisco Systems, Inc.

Distance Vector Network Discovery

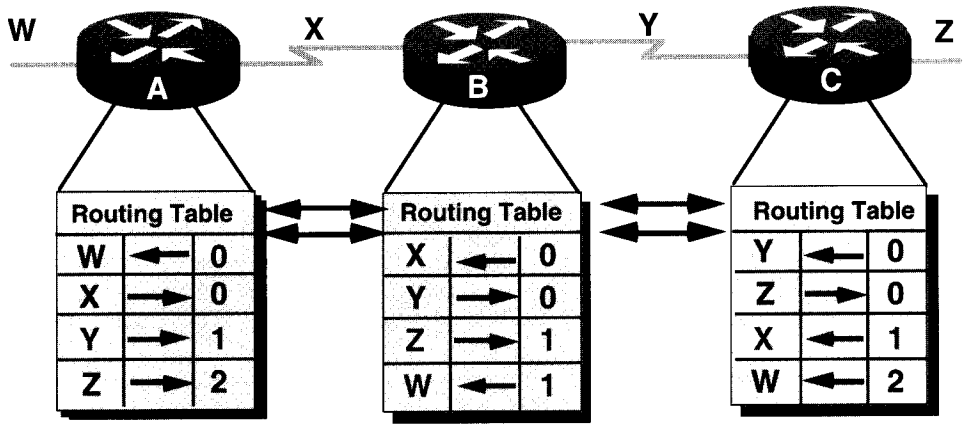

- **Routers discover the best path to destinations from each neighbor**

26

Each router using distance vector routing begins by identifying its own neighbors. In the graphic, the port to each directly connected network is shown as having a distance of 0.

As the distance vector network discovery process proceeds, routers discover the best path to destination networks based on accumulated metrics from each neighbor.

For example, router A learns about other networks based on information it receives from router B. Each of these other network entries in the routing table has an accumulated distance vector to show how far away that network is in the given direction.

Copyright © 1996, Cisco Systems, Inc.

Distance Vector Topology Changes

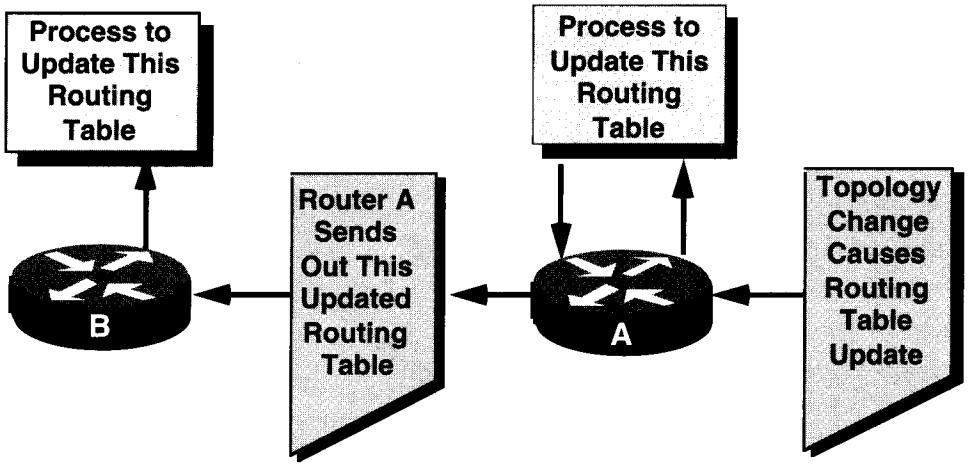

- **Updates proceed step-by-step from router to router**

27

When the topology in a distance vector protocol internetwork changes, routing table updates must occur. As with the network discovery process, topology change updates proceed step-by-step from router to router.

Distance vector algorithms call for each router to send its entire routing table to each of its adjacent neighbors. Distance vector routing tables include information about the total path cost (defined by its metric) and the logical address of the first router on the path to each network it knows about.

When a router receives an update from a neighboring router, it compares the update to its own routing table. If it learns about a better route (smaller metric) to a network from its neighbor, the router updates its own routing table. In updating its own table, the router adds the cost of reaching the neighbor router to the path cost reported by the neighbor to establish the new metric.

For example, if router B in the graphic is one unit of cost from router A, router B would add 1 to all costs reported by router A when it runs the distance vector processes to update its routing table.

Copyright © 1996, Cisco Systems, Inc.

Problem: Routing Loops

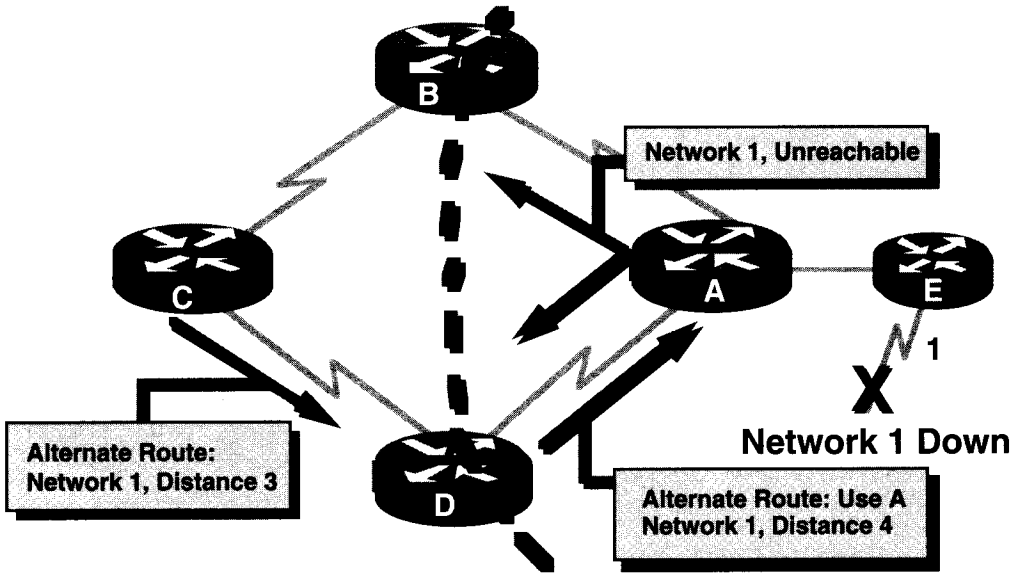

- **Alternate routes, slow convergence, inconsistent routing** 28

Routing loops can occur if the internetwork's slow convergence on a new configuration causes inconsistent routing entries. The graphic illustrates how a routing loop can occur:

- Just before the failure of network 1, all routers have consistent knowledge and correct routing tables. The network is said to have converged. Assume for the remainder of this example that router C's preferred path to network 1 is by way of router B, and router C has a distance of 3 to network 1 in its routing table.
- When network 1 fails, router E sends an update to router A. Router A stops routing packets to network 1, but routers B, C, and D continue to do so because they have not yet been informed about the failure. When router A sends out its update, routers B and D stop routing to network 1; however, router C is still not updated. To router C, network 1 is still reachable via router B. This would be the new preferred route with a metric of three hops.
- Now router C sends a periodic update to router D indicating a path to network 1 by way of router B. Router D changes its routing table to reflect this good, but erroneous, news and propagates the information to router A. Router A propagates the information to routers B and E, and so on. Any packet destined for network 1 will now loop from router C to B to A to D and back to C.

Copyright © 1996, Cisco Systems, Inc.

Problem: Counting to Infinity

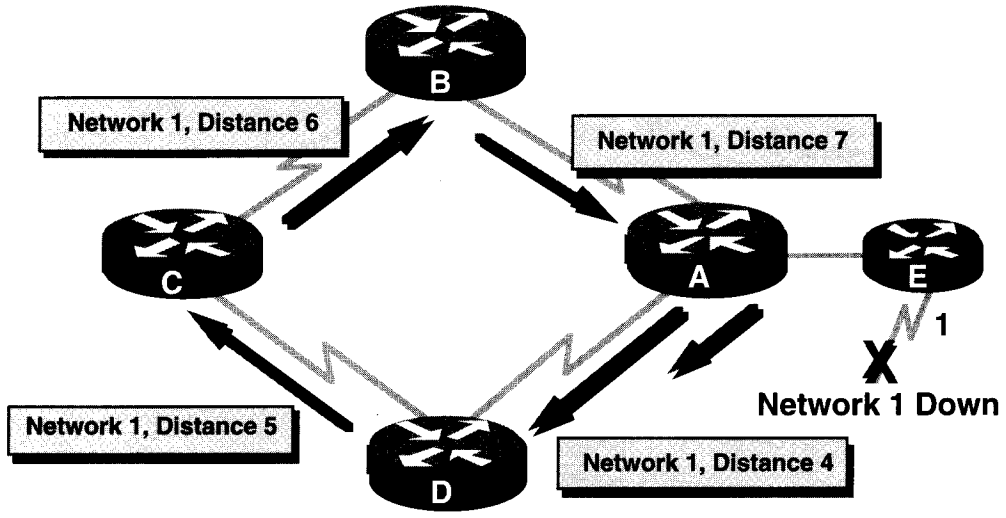

- **Routing loops increment the distance vector**

29

Continuing our example from the previous page, the invalid updates about network 1 continue to loop. Until some other process can stop the looping, the routers update each other in an inappropriate way, considering the fact that network 1 is down.

This condition, called count-to-infinity, continuously loops packets around the network, despite the fundamental fact that the destination network 1 is down. While the routers are counting to infinity, the invalid information allows a routing loop to exist.

Without countermeasures to stop the process, the distance vector of hop count increments each time the packet passes through another router. These packets loop through the network because of wrong information in the routing tables.

Copyright © 1996, Cisco Systems, Inc.

Solution: Defining a Maximum

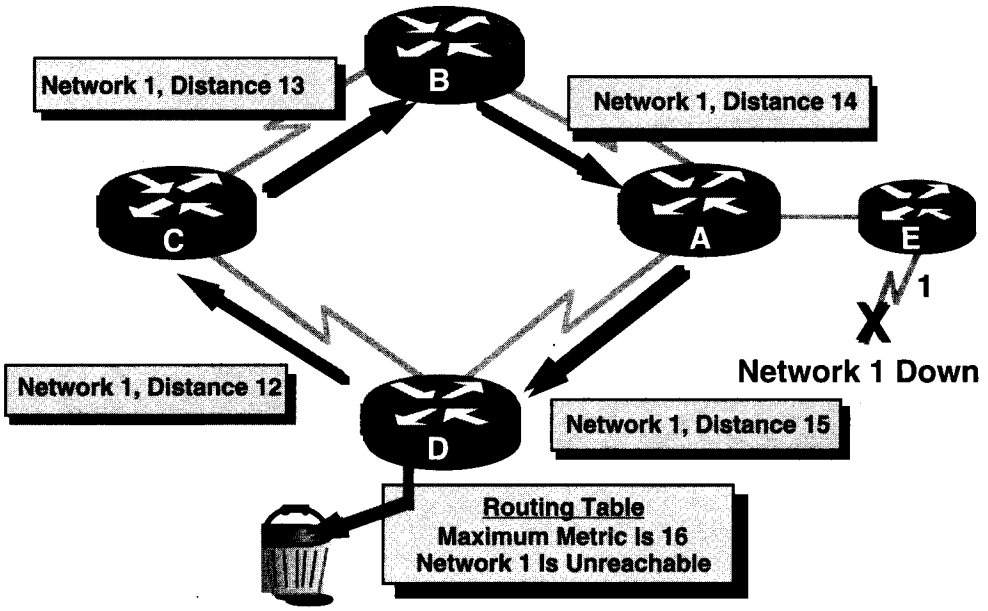

- **Specify a maximum distance vector metric as infinity** 30

Distance vector routing algorithms are self-correcting, but the routing loop problem can require a count to infinity first.

To avoid this prolonged problem, distance vector protocols define infinity as some maximum number. This number refers to a routing metric (for example, a simple hop count).

With this approach, the routing protocol permits the routing loop until the metric exceeds its maximum allowed value. The graphic shows this defined maximum as 16 hops; for hop-count distance vectors, a maximum of 15 hops is commonly used. In any case, once the metric value exceeds the maximum, network 1 is considered unreachable.

Copyright © 1996, Cisco Systems, Inc.

Solution: Split Horizon

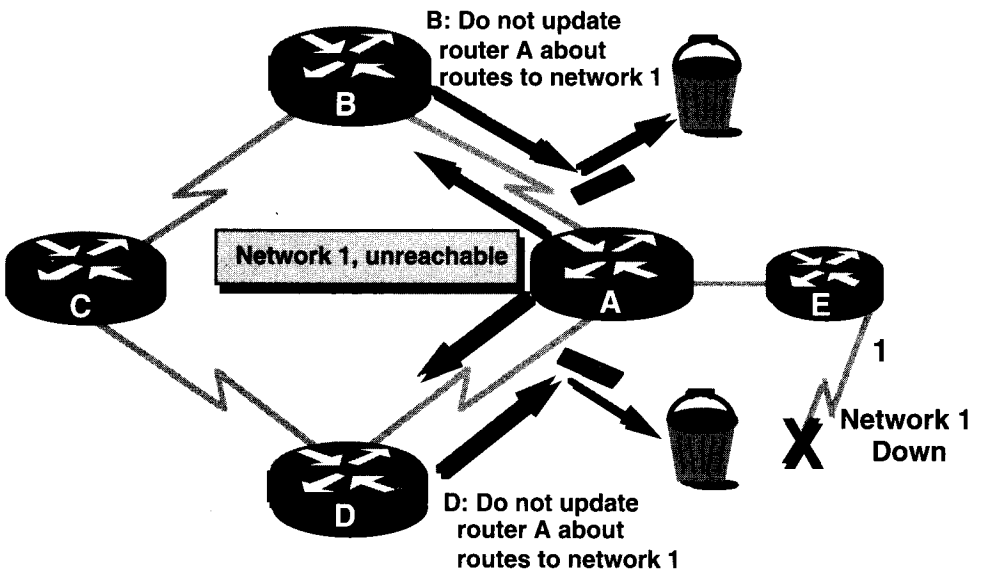

- **If you learn a protocol's route on an interface, do not send information about that route back out that interface**

31

Another possible source for a routing loop occurs when incorrect information sent back to a router contradicts the correct information it sent. Here is how this problem occurs:

- Router A passes an update to router B and router D indicating that network 1 is down. However, router C transmits an update to router B indicating that network 1 is available at a distance of 4 by way of router D. This does not violate split-horizon rules.
- Router B concludes (incorrectly) that router C still has a valid path to network 1, although at a much less favorable metric. Router B sends an update to router A advising A of the "new" route to network 1.
- Router A now determines it can send to network 1 by way of router B; router B determines it can send to network 1 by way of router C; and router C discerns it can send to network 1 by way of router D. Any packet introduced into this environment will loop between routers.

Split horizon attempts to avoid this situation. As shown in the graphic, if a table update about network 1 arrives from router A, router B or D cannot send information about network 1 back to router A. Split horizon thus reduces incorrect routing information and reduces routing overhead.

Copyright © 1996, Cisco Systems, Inc.

Solution: Route Poisoning

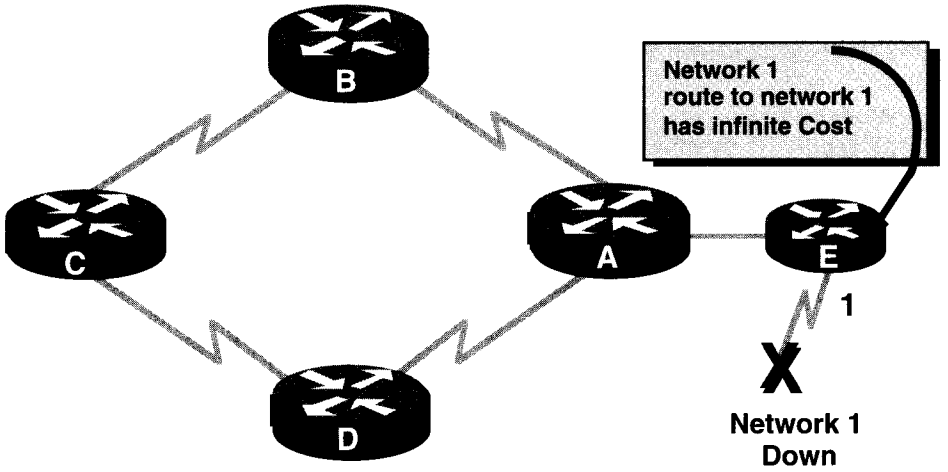

- **Router keeps an entry for the network down state, allowing time for other routers to recompute for this topology change** 32

Route poisoning offers yet another technique routers use to try to avoid the problems caused by inconsistent updates. With this technique, the router sets a table entry that keeps the network state consistent while other routers gradually converge correctly on the topology change. Used with hold-down timers, which are described on the next page, route "poisoning" is a solution to long loops.

The graphic provides the following example. When network 1 goes down, router E initiates route poisoning by entering a table entry for network 1 as having infinite cost (that is, being unreachable). By poisoning its route to network 1, router E is not susceptible to other incorrect updates about network 1 coming from neighboring routers that might claim to have a valid alternate path. This can work with the hold-down mechanism described on the next page.

Router E keeps this poison-route entry for several update cycles. The poisoned router can trigger an update about network 1 in neighbor routers (as well as the other routers in the internetwork).

Route poisoning and triggered updates speed up convergence because the routers do not have to wait for update intervals before advertising the poisoned route. This can hasten the spread of updated path information about network 1 as these other routers recompute their distance vector tables and converge on the topology change.

Copyright © 1996, Cisco Systems, Inc.

Solution: Hold-Down Timers

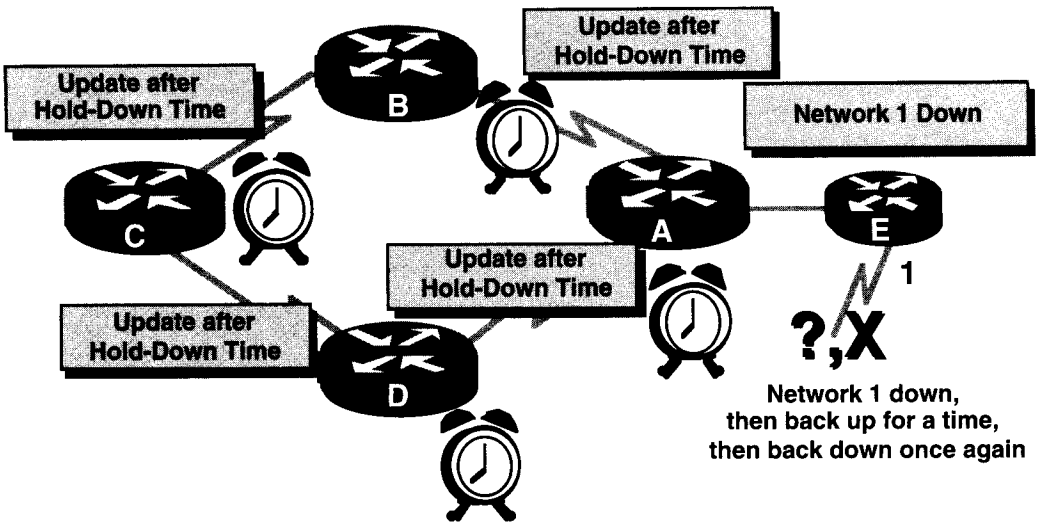

- **Routers ignore network update information for some period**

33

You can avoid the count-to-infinity problem by using hold-down timers, which work as follows:

- When a router receives an update from a neighbor indicating that a previously accessible network is now inaccessible, the router marks the route as inaccessible and starts a hold-down timer. If at any time before the hold-down timer expires an update is received from the same neighbor indicating that the network is again accessible, the router marks the network as accessible and removes the hold-down timer.
- If an update arrives from a different neighboring router with a better metric than originally recorded for the network, the router marks the network as accessible and removes the hold-down timer.
- If at any time before the hold-down timer expires an update is received from a different neighboring router with a poorer metric, the update is ignored. Ignoring an update with a poorer metric when a hold-down is in effect allows more time for the knowledge of a disruptive change to propagate through the entire network.

Hold-down timers work with route poisoning.

Copyright © 1996, Cisco Systems, Inc.

Link-State Concept

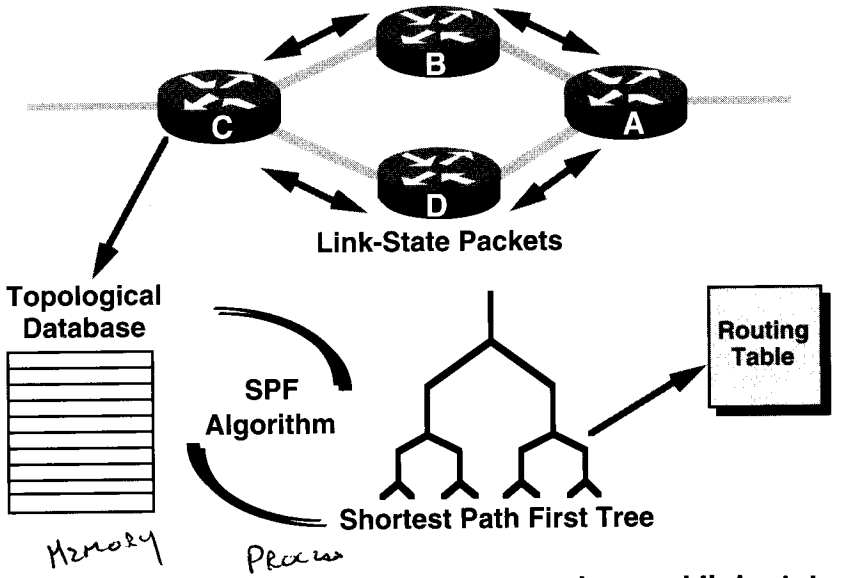

- **After initial flood, pass small event-triggered link-state updates to all other routers**

34

The second basic algorithm used for routing is the link-state algorithm.

Link-state-based routing algorithms—also known as shortest path first (SPF) algorithms—maintain a complex database of topology information. Whereas the distance vector algorithm has nonspecific information about distant networks and no knowledge of distant routers, a link-state routing algorithm maintains full knowledge of distant routers and how they interconnect.

Link-state routing uses link-state packets (LSPs), a topological database, the SPF algorithm, the resulting SPF tree, and finally, a routing table of paths and ports to each network. The following pages cover these processes and databases in more detail.

Engineers have implemented this link-state concept in Open Shortest Path First (OSPF) routing. RFC 1583 contains a description of OSPF link-state concepts and operations.

LINK STATE REQUIRES ADDITIONAL MEMORY & PROCESSING POWER.

OSPF - IP
NLSP - IPX
IS-IS - IP, DEC
DEC PHASE V - DEC

DUAL ALGORITHM
EIGRP - IP, IPX, AT

Copyright © 1996, Cisco Systems, Inc.

Link-State Network Discovery

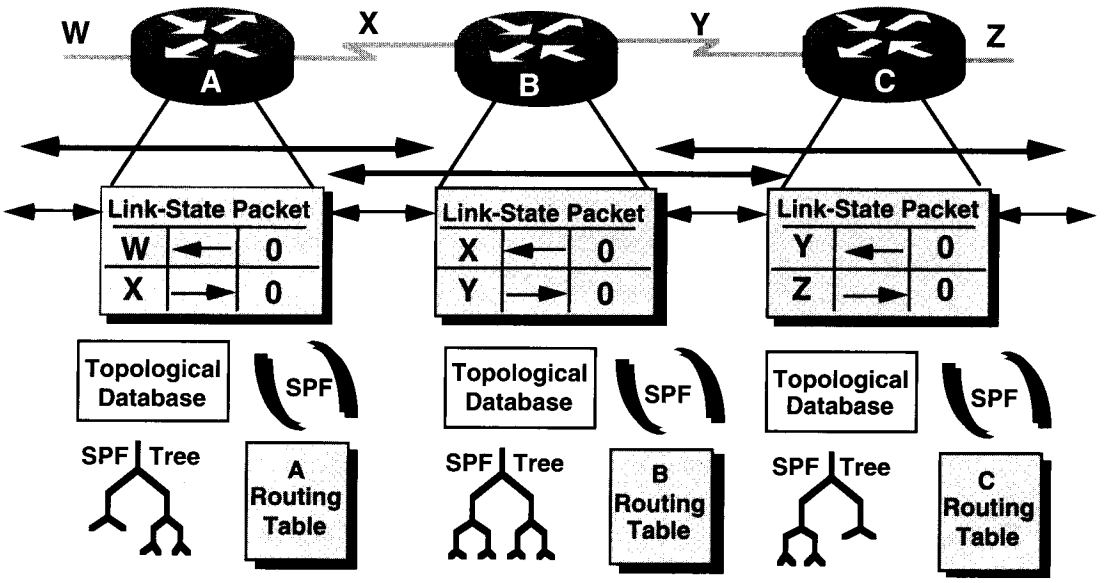

- **Routers calculate the shortest path to destinations in parallel**

35

Network discovery for link-state routing uses the following processes:

- Routers exchange LSPs with each other. Each router begins with directly connected networks for which it has direct link-state information.
- Next, each router in parallel with one another constructs a topological database consisting of all the LSPs from the internetwork.
- The SPF algorithm computes network reachability, determining the shortest path first to each other network in the link-state protocol internetwork. The router constructs this logical topology of shortest paths as an SPF tree. With itself as root, this tree expresses paths from the router to all destinations.
- The router lists its best paths and the ports to these destination networks in the routing table. It also maintains other databases of topology elements and status details.

After the routers dynamically discover the details of their internetwork, they can use the routing table for switching packet traffic.

Copyright © 1996, Cisco Systems, Inc.

Link-State Topology Changes

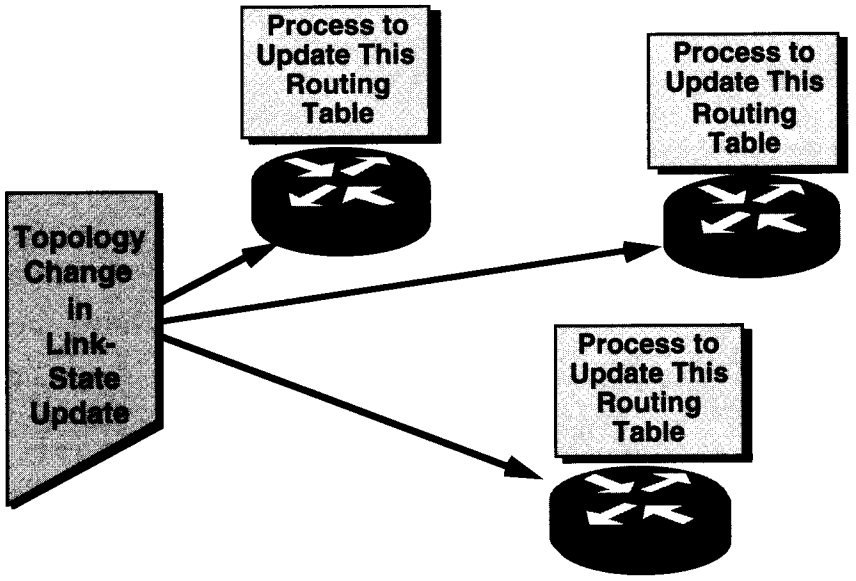

- **Update processes proceed using the same link-state update**

36

Link-state algorithms rely on using the same link-state updates. Whenever a link-state topology changes, the routers that first become aware of the change send information to other routers or to a designated router that all other routers can use for updates. This entails the propagation of common routing information to all routers in the internetwork. To achieve convergence, each router does the following:

- Keeps track of its neighbors: the neighbor's name, whether the neighbor is up or down, and the cost of the link to the neighbor.
- Constructs an LSP that lists its neighbor router names and link costs. This includes new neighbors, changes in link costs, and links to neighbors that have gone down.
- Sends out this LSP so that all other routers receive it.
- When it receives an LSP, records the LSP in its database so that it can store the most recently generated LSP from each other router.
- Using accumulated LSP data to construct a complete map of the internetwork topology, proceeds from this common starting point to rerun the SPF algorithm and compute routes to every network destination.

Each time an LSP causes a change to the link-state database, the link-state algorithm recalculates the best paths and updates the routing table. Then every router takes the topology change into account as it determines the shortest paths to use for packet switching.

Copyright © 1996, Cisco Systems, Inc.

Link-State Concerns

- **Processing and memory required for link-state routing**

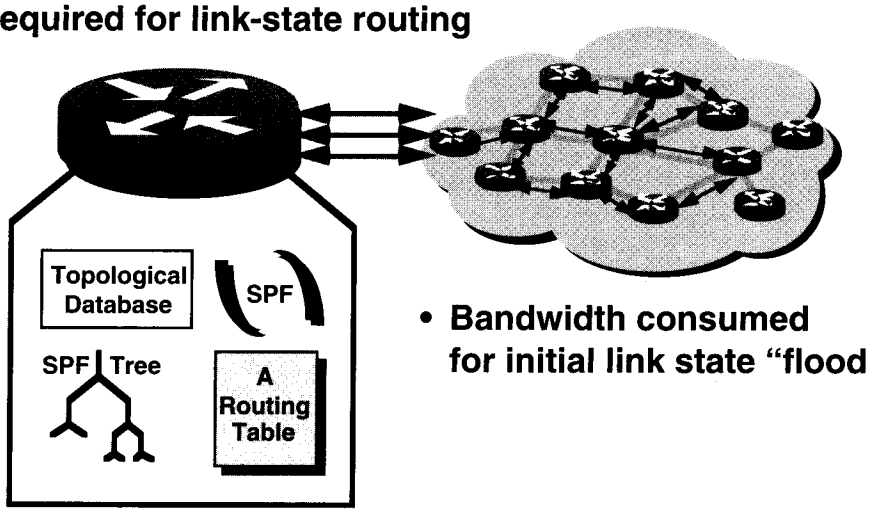

- **Bandwidth consumed for initial link state "flood"**

37

There are two link-state concerns:

- Processing and memory requirements—Running link-state routing protocols in most situations requires that routers use more memory and perform more processing. Network administrators must ensure that the routers they select are capable of providing these resources for routing.

  Routers keep track of their neighbors and the networks they reach through other routing nodes. For link-state routing, memory must hold information from various databases, the topology tree, and the routing table.

  Computing the shortest path first with Dijkstra's algorithm requires a processing task proportional to the number of links in the internetwork times the number of routers in the network.

- Bandwidth requirements—Another cause for concern involves the bandwidth consumed for initial link-state packet flooding. During the initial discovery process, all routers using link-state routing protocols send LSPs to all other routers. This action floods the internetwork as routers make their peak demand for bandwidth, and temporarily reduces the bandwidth available for routed traffic that carries user data.

  After this initial flooding, link-state routing protocols generally require only internetwork bandwidth to send infrequent or event-triggered LSPs that reflect topology changes.

Copyright © 1996, Cisco Systems, Inc.

Problem: Link-State Updates

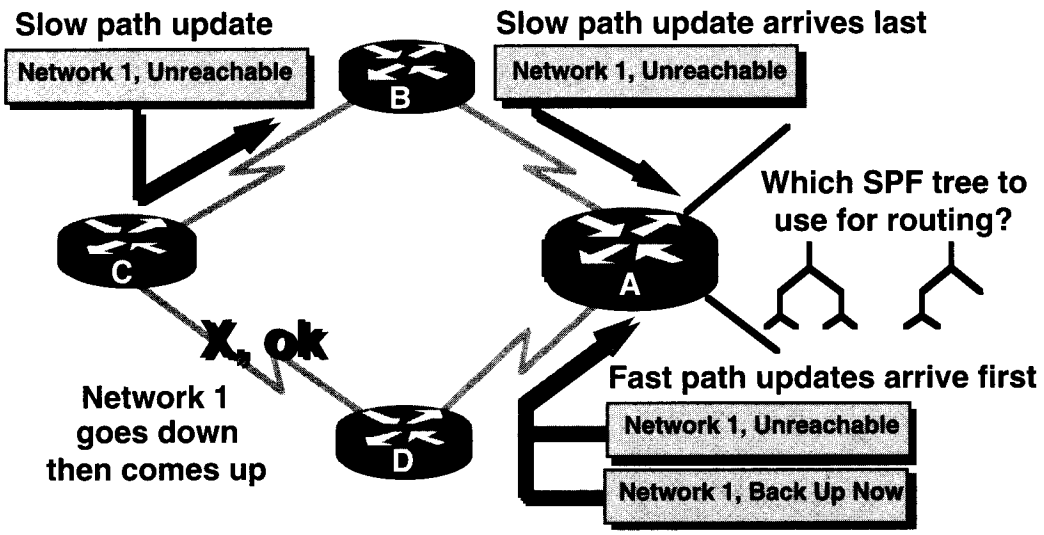

- **Unsynchronized updates, inconsistent path decisions**

38

The most complex and critical aspect of link-state routing is making sure that all routers get all the LSPs necessary. Routers with different sets of LSPs will calculate routes based on different topological data. Then routes become unreachable as a result of the disagreement among routers about a link. Here is an example of inconsistent path information:

- Suppose that network 1 between routers C and D goes down. As discussed earlier, both routers construct an LSP to reflect this unreachable status.
- Soon afterward, network 1 comes back up; another LSP reflecting this next topology change is needed.
- If the original "Network 1, Unreachable" message from router C uses a slow path for its update, that update comes later. This LSP can arrive at router A after router D's "Network 1, Back Up Now" LSP.
- With unsynchronized LSPs, router A can face a dilemma about which SPF tree to construct: Does it use paths containing network 1 or without network 1, which was most recently reported as unreachable?

If LSP distribution to all routers is not done correctly, link-state routing can result in invalid routes.

Copyright © 1996, Cisco Systems, Inc.

Link-State Update Problems (cont.)

- **Synchronizing large networks—which network topology updates are correct?**
- **Router startup—order of start alters the topology learned**
- **Partitioned regions—slow updating part separated from fast updating part**

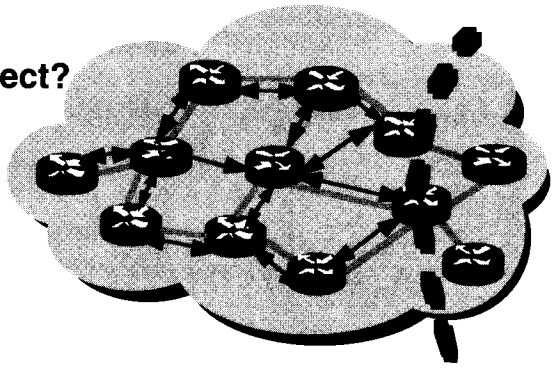

39

Scaling up with link-state protocols on very large internetworks can intensify the problem of faulty LSP distribution.

If one part of the internetwork comes up first with other parts coming up later, the order for sending and receiving LSPs will vary. This variation can alter and impair convergence. Routers might learn about different versions of the topology before they construct their SPF trees and routing tables.

On a large internetwork, parts that update more quickly can cause problems for parts that update more slowly. Routers sending out LSPs cannot assume they will be correctly transported by following existing routing table entries because these entries might not reflect the current topology.

With faulty updates, LSPs can multiply as they propagate through the internetwork, unproductively consuming more and more bandwidth.

Eventually a partition can split the internetwork into a fast updating part and a slow updating part. Then network administrators must troubleshoot the link-state complexities to restore acceptable connectivity.

Copyright © 1996, Cisco Systems, Inc.

Solution: Link-State Mechanisms

- **Reduce the need for resources**
  - "Dampen" update frequency
  - Target link-state updates to multicast
  - Use link-state area hierarchy for topology
  - Exchange route summaries at area borders
- **Coordinate link-state updates**
  - Use time stamps
  - Update numbering and counters
  - Manage partitioning using an area hierarchy

40

Link-state routing has several techniques for preventing or correcting potential problems arising from resource requirements and LSP distribution.

- A network administrator can reduce the periodic distribution of LSPs so that updates only occur after some long, configurable duration. The dampening does not interfere with LSP updates triggered by topology changes.
- LSP updates can go to a multicast group rather than in a flood to all routers. On interconnected LANs, you can use one or more designated routers as the target depository for LSP transmissions. Other routers can use their designated routers as a specialized source of consistent topology data.
- In large networks you can set up a hierarchy made up of different router levels. A router in one area of the hierarchical domains does not need to store and process LSPs from other routers not located in its area.
- For problems of LSP coordination, link-state implementations can allow for LSP time stamps, sequence numbers, aging schemes, and other related mechanisms to help avoid inaccurate LSP distribution or uncoordinated updates.
- The partitioning of an internetwork can be actively managed with a hierarchy if the link-state protocol provides for hierarchical management. Then routers can concentrate on only the routers within their domain or area, and they can depend on special routers at the domain borders for external routing information.

Copyright © 1996, Cisco Systems, Inc.

Comparing Distance Vector Routing to Link-State Routing

Distance Vector	Link-State
Views net topology from neighbor's perspective	Gets common view of entire network topology
Adds distance vectors from router to router	Calculates the shortest path to other routers
Frequent, periodic updates: slow convergence	Event-triggered updates: faster convergence
Passes copies of routing table to neighbor routers	Passes link-state routing updates to other routers

41

You can compare distance-vector routing to link-state routing in several key areas:

- Distance vector routing gets all topological data from the perspective it receives from processing the routing table information of its neighbors. Link-state routing obtains a wide view of the entire internetwork topology by accumulating all necessary LSPs.
- Distance vector routing determines the best path by adding to the metric value it receives as tables move from router to router. For link-state routing, each router works in parallel to calculate its own shortest path to destinations.
- With most distance vector routing protocols, updates for topology changes come in periodic table updates. These tables pass incrementally from router to router, usually resulting in slower convergence.
- With link-state routing protocols, updates are usually triggered by topology changes. Relatively small LSPs passed to all other routers, or a multicast group of routers, usually result in faster time to converge on any internetwork topology change.

RIP - IP (HOPS) IP X (TICKS / HOPS) XNS (HOP)

IGRP - IP (BANDWIDTH)

RTMP AT (HOPS)

RTP VINES (COST)

DEC PHV DSE (COST)

OSPF IP (COST)

NLSP IPX

EIGRP (4 BANDWIDTH)

Copyright © 1996, Cisco Systems, Inc.

What Is Best? It Depends

Issue	Concern	Example Questions
Technical	Performance to meet specific needs	Metrics adequate for network size? Any load sharing?
Business	Conformity with enterprise policies and priorities	Proven technology? Multivendor support? Standards based?
Operational	Simplicity of network setup and management	Easy to configure? Able to handle several routed protocols?

- **With routing protocols, no one type fits all networks**

42

Routing decisions are not always based on the quickest, shortest, cheapest, or most reliable path. The answers to your technical concerns and questions might provide only some of the factors that determine your choice.

Business concerns can also influence routing policy. Conformance with the policies, priorities, and partnerships of an organization impacts routing choices. For example, one routing selection might be considered more desirable because it uses the facilities of a partner or avoids the facilities of a competitor. Multivendor support or standards conformity might outweigh technical superiority.

Operational issues such as the concern for network simplicity are also important. For the chosen routing protocol to properly fit some organizations, it must be easy to set up and manage. It must handle several routed protocols without requiring multiple inconsistent and complex configuration templates. Also, network administrators can sustain their careers by using proven technologies and thereby avoiding risk.

Size of Network

Diversity of Router Vendors.

Copyright © 1996, Cisco Systems, Inc.

Hybrid Routing

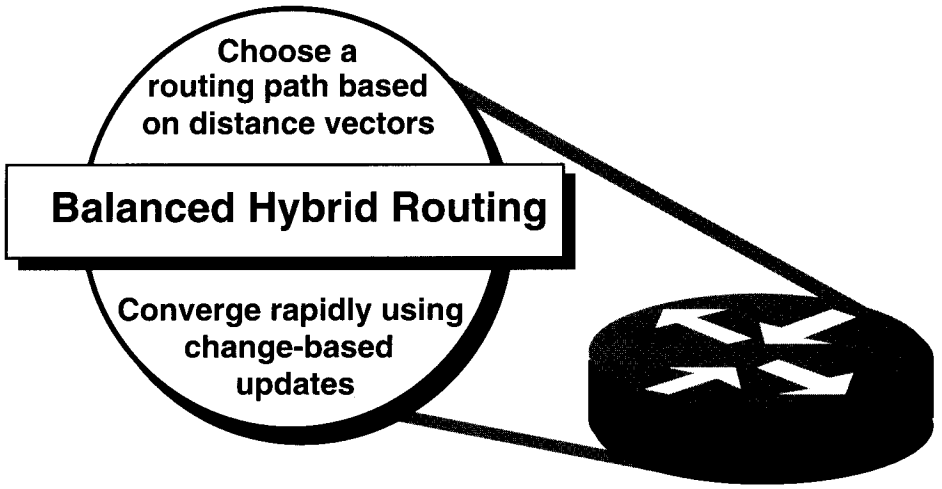

- **Share attributes of both distance-vector and link-state routing**

43

This chapter so far has presented the two major types of routing protocols: distance vector and link-state.

An emerging third type of routing protocol combines aspects of both. This third type is called balanced hybrid in this course.

The balanced hybrid routing protocol uses distance vectors with more accurate metrics to determine the best paths to destination networks. However, it differs from most distance vector protocols by using topology changes to trigger routing database updates.

The balanced hybrid routing type converges more rapidly, like the link-state protocols. However, it differs from these protocols by emphasizing economy in the use of required resources such as bandwidth, memory, and processor overhead.

Examples of balanced hybrid protocols are OSI's Intermediate System-to-Intermediate System (IS-IS) routing and Cisco's Enhanced Interior Gateway Routing Protocol (Enhanced IGRP).

Copyright © 1996, Cisco Systems, Inc.

Integrated Routing

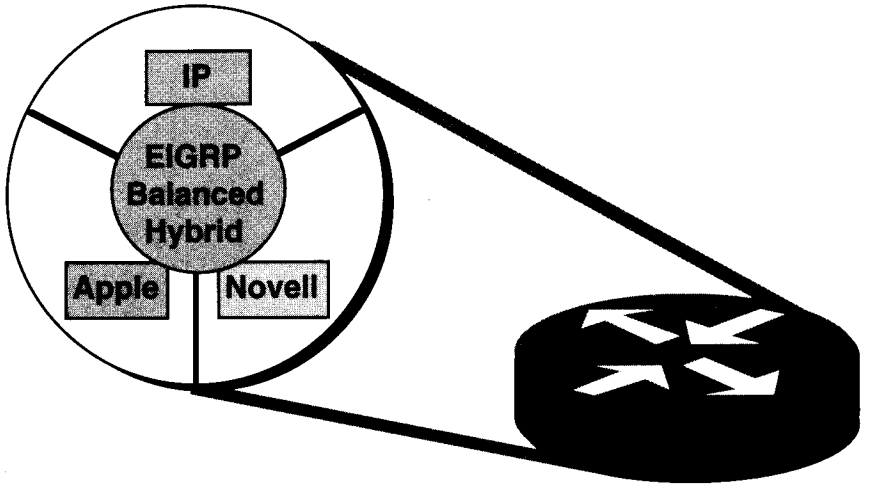

- **Use routing decisions to improve paths for routed protocols**

44

Earlier material in this chapter described multiprotocol routing operating as ships-in-the-night routing—routing with no protocol directly affecting any other protocol.

An alternative form of multiprotocol routing uses a single integrated routing algorithm that does the following:

- Supports path selection and packet switching for more than one routed protocol.
- Carries routing updates that are usable by several routed protocols.
- Replaces the native routing algorithm, which is no longer required.
- Creates separate routing tables for each of the routed protocols.
- Saves network and router resources.
- Simplifies the administrator's operational tasks.

An example of an integrated routing protocols is Enhanced IGRP, which is a Cisco proprietary routing protocol that integrates support for IP, AppleTalk, and Novell IPX. Enhanced IGRP uses a distance vector algorithm based on Cisco's IGRP.

Copyright © 1996, Cisco Systems, Inc.

LAN-to-LAN Routing

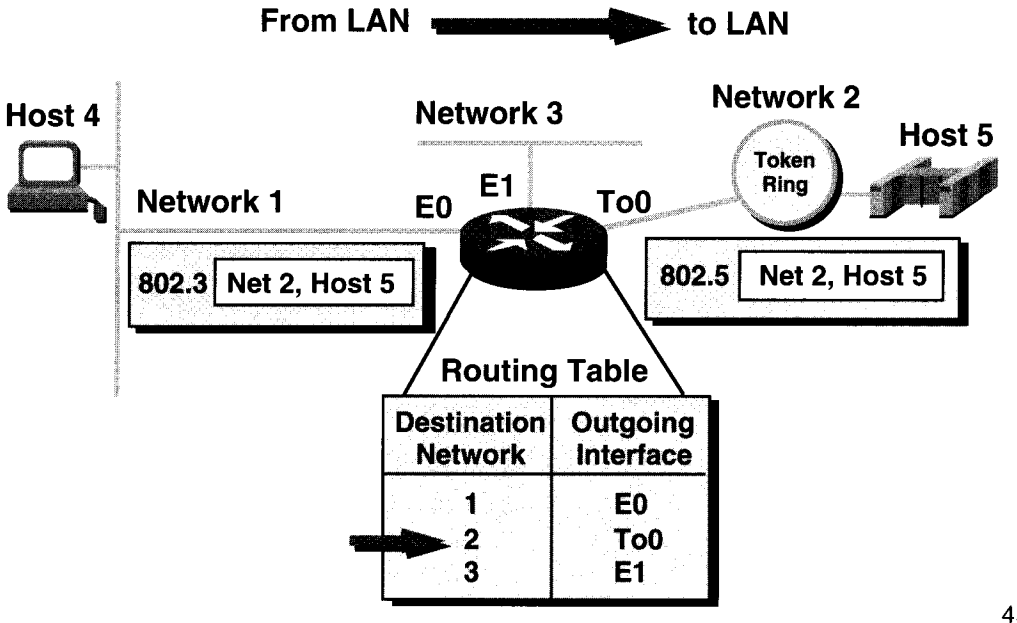

45

The network layer must relate to and interface with various lower layers. Routers must be capable of seamlessly handling packets encapsulated into different lower-level frames without changing the packets' Layer 3 addressing.

The graphic shows an example of this with LAN-to-LAN routing. In this example, packet traffic from source host 4 on Ethernet network 1 needs a path to destination host 5 on network 2. The LAN hosts depend on the router and its consistent network addressing to find the best path.

When the router checks its router table entries, it discovers that the best path to destination network 2 uses outgoing port To0, the interface to a Token Ring LAN.

Although the lower-layer framing must change as the router switches packet traffic from the Ethernet on network 1 to the Token Ring on network 2, the Layer 3 addressing for source and destination remains the same. In the graphic, the destination address remains Net 2, Host 5 despite the different lower-layer encapsulations.

Copyright © 1996, Cisco Systems, Inc.

LAN-to-WAN Routing

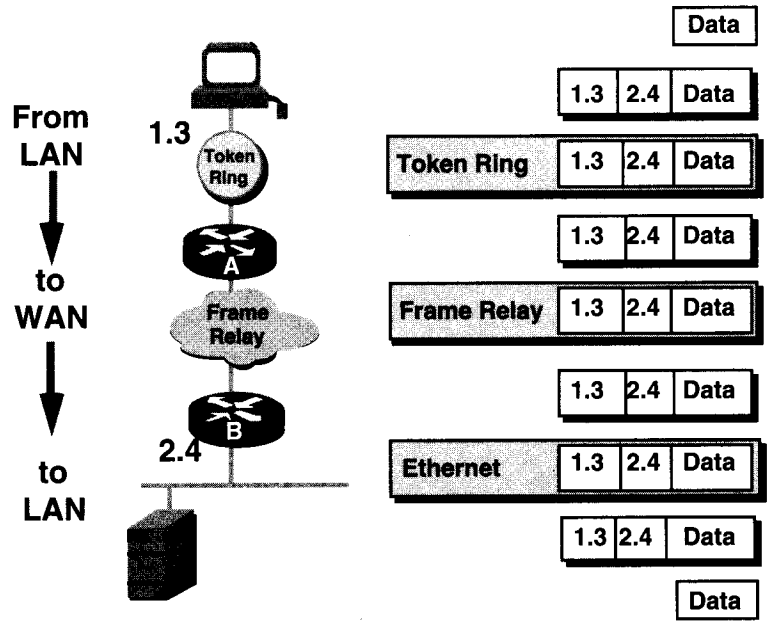

46

The network layer must relate to and interface with various lower layers for LAN-to-WAN traffic. As an internetwork grows, the path taken by a packet might encounter several relay points and a variety of data-link types beyond the LANs. For example, in the graphic, a packet from the top workstation at address 1.3 must traverse three data links to reach the file server at address 2.4 shown on the bottom.

- The workstation sends a packet to the file server by encapsulating the packet in a Token Ring frame addressed to router A.
- When router A receives the frame, it removes the packet from the Token Ring frame, encapsulates it in a Frame Relay frame, and forwards the frame to router B.
- Router B removes the packet from the Frame Relay frame and forwards the packet to the file server in a newly created Ethernet frame.
- When the file server at 2.4 receives the Ethernet frame, it extracts and passes the packet to the appropriate upper-layer process.

The routers enable LAN-to-WAN packet flow by keeping the end-to-end source and destination addresses constant while encapsulating the packet at the port to a data link that is appropriate for the next hop along the path.

Copyright © 1996, Cisco Systems, Inc.

Layer Decapsulation

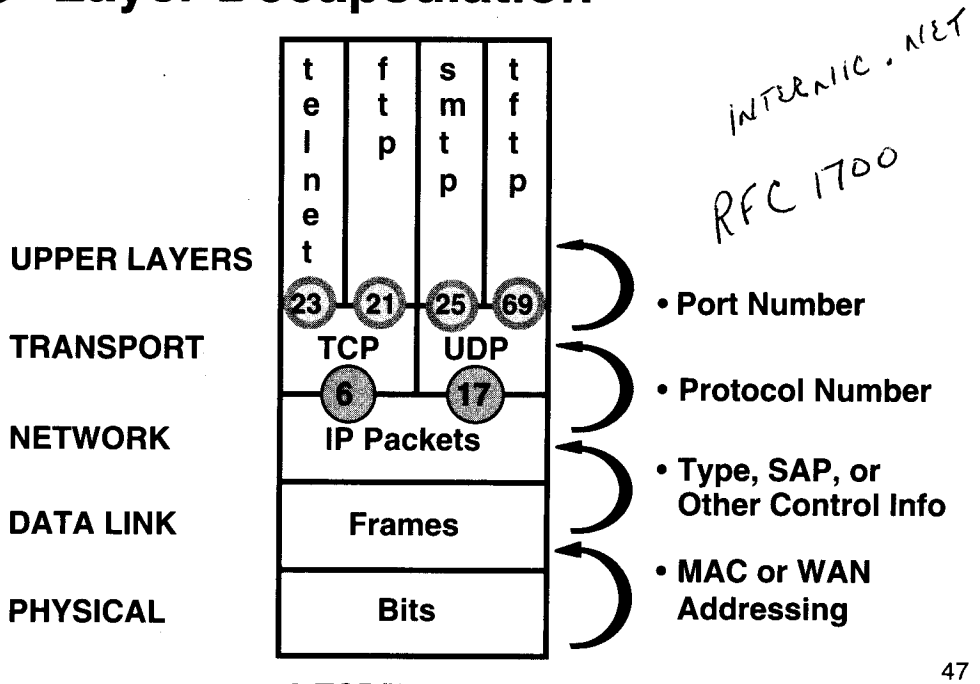

Having touched on each of the major ISO/OSI reference model layers, you can tie them all together with mechanisms like those shown in the graphic for TCP/IP. These transition mechanisms occur between layers. Information encapsulates as it makes the transition into the next lower layer, ultimately sending bits out a physical interface with clocking. The graphic illustrates decapsulation—data making the transition from lower to upper layers. Information decapsulates as it strips away lower-layer headers and data-link trailers.

- Bits entering from an interface assemble into frames. Among the fields in the frames will be MAC addresses if on a LAN, or WAN addresses appropriate for the given WAN service.
- In the frame, a type field (Ethernet II), SAP (LLC), or other control information field carries information about the packet contained as data within the frame.
- IP packets contain a field in the header to indicate the protocol of the segment carried in its data. This number indicates whether the Layer 4 transport protocol will be TCP (6) or UDP (17).
- At the transport layer, the segment includes a port number to indicate the upper-layer protocol type carried in its data. For example, a TCP segment carries port 23 as a well-known port for Telnet.
- The top layer provide network applications. These can provide network functions for end users directly (as with FTP), or can work through other applications (as SMTP provides a gateway between different e-mail applications).

Copyright © 1996, Cisco Systems, Inc.

Cisco Router Configuration

Understand Multiple Protocols

Interconnect Multiple Media

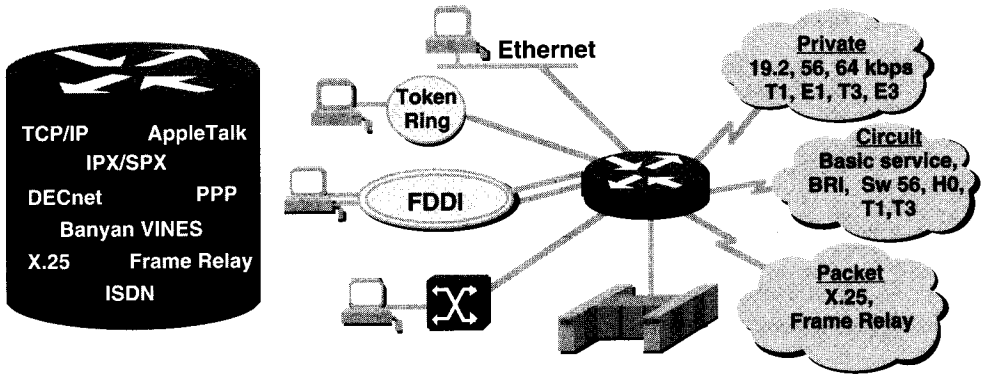

- **Your tasks for the remainder of this course**

48

Routers are devices that implement the network service. They provide interfaces for a wide range of links and subnetworks at a wide range of speeds. Routers are active and intelligent network nodes and thus can participate in managing the network. Routers manage networks by providing dynamic control over resources and supporting the tasks and goals for internetworks: connectivity, reliable performance, management control, and flexibility.

In addition to the basic switching and routing functions, routers have implemented a variety of value-added features that help to improve the cost-effectiveness of the internetwork. These features include sequencing traffic based on priority and traffic filtering.

Typically, routers are required to support multiple protocol stacks, each with its own routing protocols, and to allow these different environments to operate in parallel. In practice, routers also incorporate bridging functions and can serve as a limited form of hub.

This course is about the fundamentals of configuring software on Cisco routers. In the modules that follow, you will learn operations and techniques for configuring the Cisco routers in class to operate the protocols and many of the media shown in the graphic.

Copyright © 1996, Cisco Systems, Inc.

Summary

Internetworking functions of the network layer include network addressing and best path selection for traffic

Network addressing uses one part to identify the path used by the router and one part for ports or devices on the net

Rout*ed* protocols direct user traffic, while rout*ing* protocols work between routers to maintain path tables

Network discovery for distance vector involves exchange of routing tables; problems can include slower convergence

For link-state, routers calculate the shortest paths to other routers; problems can include inconsistent updates

Balanced hybrid routing uses attributes of both link-state and distance vector, applying paths to several protocols

49

Copyright © 1996, Cisco Systems, Inc.

Exercise: Network Layer and Path Determination

Problem 1

Objective: Contrast the network discovery and update processes in distance vector routing with those in link-state routing.

Fill in the blank lines in the graphic below with the letter of the statement that best describes that aspect of the routing protocol type. For example, enter the letter C in the Discovery box for Distance Vector Routing if you think that distance vector routing develops a view from neighbor routers' perspectives.

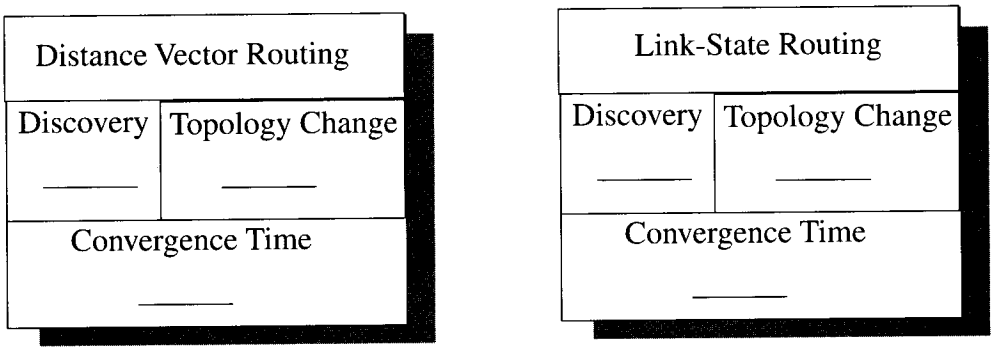

Note You will not find a place for every statement.

A) Starts from a common view of internetwork topology.

B) Uses periodic updates, resulting in relatively slow convergence.

C) Develops a view from neighbor routers' perspectives.

D) Provides end-to-end load sharing.

E) Passes an updated routing table from neighbor to neighbor.

F) Events trigger updates for relatively fast convergence.

G) Processes updates in parallel with other routers.

H) Uses distance vector metrics and change-based updates.

Copyright © 1996, Cisco Systems, Inc.

Problem 2

Objective: List problems that each routing type encounters when dealing with topology changes, and describe techniques to reduce the number of these problems.

List one routing problem and indicate what type of routing is most likely to have the problem. Then list at least two techniques to solve (or reduce) that problem.

Problem 3

Objective: Explain the services of separate and integrated multiprotocol routing.

Describe at least one key difference between a separate multiprotocol routing environment and an integrated multiprotocol environment.

Problem 4

In the graphic, the rectangle identified as circle A shows a simplified Ethernet II frame. Other rectangles inside this frame represent encapsulations from various layers of TCP/IP. Identify these encapsulated components. Write the name for each component on the line pointing to the circled letter of the component. For example, the correct answer for the Ethernet II frame is provided on the line for component A.

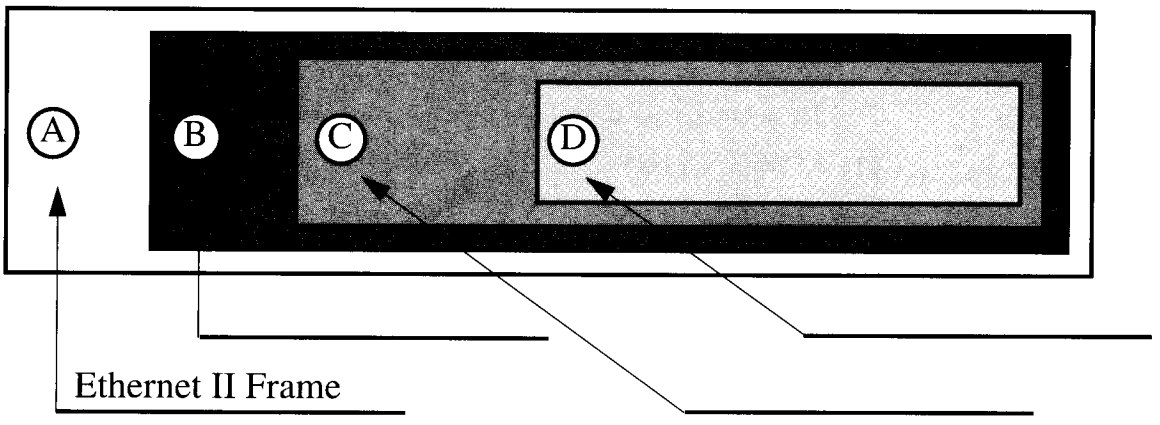

Copyright © 1996, Cisco Systems, Inc.

Problem 5

Refer to the diagram shown in the problem 4. The following information represents part of the contents possible for TCP/IP layers. Each item on the list refers to a layer transition as information is decapsulated:

1. The packet carries IP.
2. Use Trivial File Transfer Protocol (TFTP).
3. The program begins using the port number 20.
4. Use the IP protocol number 17 (User Datagram Protocol).
5. The destination interface is 00000c123456; the source is 00000c135791.

On the following space provided, match the number of the contents description with its component letter. For example, write the answer 2 next to the letter D if *Use Trivial File Transfer Protocol (TFTP)* is your answer for layer transition information that component D could contain.

Note The numbered list of content descriptions intentionally contains an extra statement.

A) _____

B) _____

C) _____

D) _____

Problem 6

A) EIGRP and ISIS are examples of a new approach called

_____.

B) Must the router be configured with network-layer addresses for most protocols?

[ ] Yes [ ] No

C) Does the network-layer logical address change over each link?

[ ] Yes [ ] No

D) What does ARP stand for and what does it do?

_____

_____

E) Is a frame routable?

[ ] Yes [ ] No

Copyright © 1996, Cisco Systems, Inc.

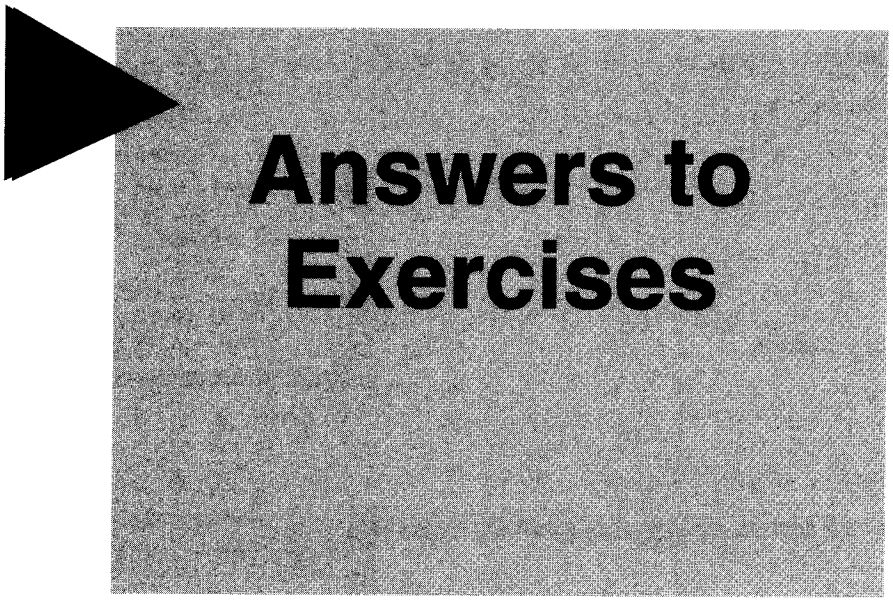

53

Answers to Exercises

Copyright © 1996, Cisco Systems, Inc.

Exercise: Network Layer Basics

Problem 1

Establishes network addresses.

Selects the best path through an internetwork.

Uses a routing protocol between routers.

Uses a routed protocol to carry user packets.

Uses a two-part address.

Sets up and maintains routing tables.

Discovers networks.

Adapts to internetwork topology changes.

Contains broadcasts.

Problem 2

The two general parts of a Layer 3 address are a network part and a node part.

Problem 3

A) 131.108.3.1 is an IP address; 131.108 is the network number, 3.1 is the host number.

B) 1000.128 is an AppleTalk DDP address; 1000 is the network number, and 128 is the host number.

C) abadcafe.0000.0c56.de33 is a Novell IPX address; abadcafe is the network number, and 0000.0c56.de33 is the node number (adapted from a MAC address).

D) 31060004085551 is an X.121 address for X.25; 3106 is the DNIC, and 0004085551 is the NTN.

Problem 4

A) Routing table.

B) Routing protocol.

C) One for each network-layer protocol. Packets are switched independently. This approach is sometimes described as "ships-in-the-night" routing.

Copyright © 1996, Cisco Systems, Inc.

Exercise: Network Layer and Path Determination

Problem 1

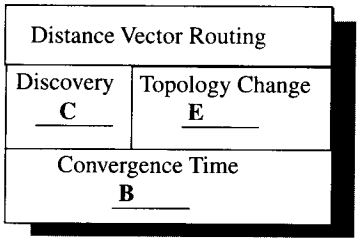

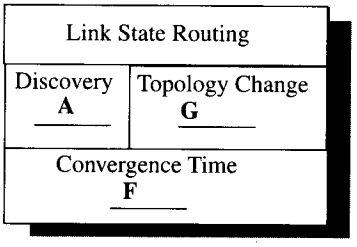

Problem 2

Distance vector problems—Routing loops, counting to infinity, slow convergence.

Techniques to solve these problems include defining infinity as a maximum, split horizon, route poisoning, and hold-down timers.

Link-state problems—Processor and memory requirements, initial flood on bandwidth, faulty LSP distribution, unsynchronized updates, and partitioned regions.

Techniques to solve these problems include multicast updates; dampened update frequency; area hierarchy; route summaries at borders; and time stamps, sequence numbering, counters, and designated routers for the LSPs.

Problem 3

In a separate multiprotocol routing environment, the several configured protocols operate like ships in the night.

With an integrated multiprotocol routing environment, the several configured protocols share the results of the integrated routing algorithm.

Copyright © 1996, Cisco Systems, Inc.

Problem 4

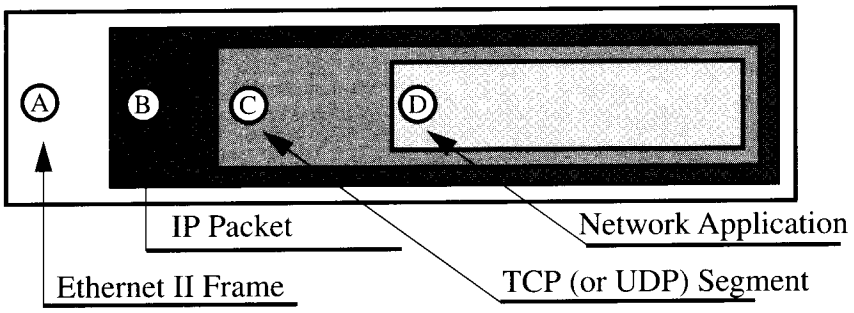

Problem 5

A) 1 (In the frames type field.)

B) 4 (In the protocol field of the datagram.

C) 3 (The well-known port number for FTP data.)

D) 2 (The example provided in the problem instructions is correct.)

Problem 6

A) Integrated routing.

B) Yes.

C) No.

D) Address Resolution Protocol. Translates network (logical) addresses to physical (MAC) addresses.

E) No. A frame is a data-link entity—it is not routable.

Copyright © 1996, Cisco Systems, Inc.