

Cisco Technical Expert



Cisco Router *Handbook*

- DETAILED COVERAGE OF INTERNETWORKING LEGACY SYSTEMS
- PRACTICAL INSIGHT INTO ALL ROUTING PROTOCOLS, INCLUDING RIP, EIGRP, OSPF, AND BGP
- EXTENSIVE DISCUSSION OF INTERNETWORKING DESIGN

GEORGE C. SACKETT

CISCO ROUTER HANDBOOK

“This reference manual will serve as an excellent guide for any field engineer looking for a concise answer to almost any Cisco-related problem. For the mid-level to senior engineer this book will be an excellent research and learning tool when considering implementation of a new Cisco-based technology solution.”

—*E. Gary Hauser Jr. Senior Instructor/Consultant,
CCIE #4489, CCSI#99166
Chesapeake Computer Consultants, Inc.*

“The one reference you take into the field. A compact source of solutions to the most common Cisco design and configuration questions. The first place to look when researching your networking projects.”

—*Jim Bender, Internetworking Consultant, RPM Consulting*

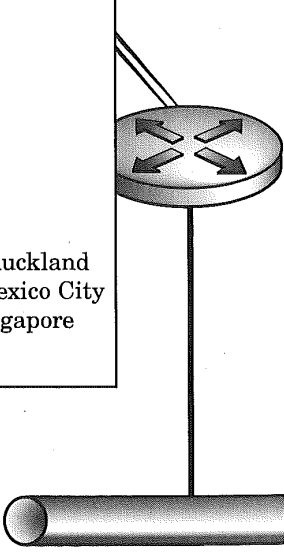



Cisco Router Handbook

George C. Sackett

McGraw-Hill

New York • San Francisco • Washington, D.C. • Auckland
Bogotá • Caracas • Lisbon • London • Madrid • Mexico City
Milan • Montreal • New Delhi • San Juan • Singapore
Sydney • Tokyo • Toronto



Library of Congress Cataloging-in-Publication Data

Sackett, George C.

The Cisco router handbook George Sackett.

p. cm.

Includes index.

ISBN 0-07-058098-7

1. Routers (Computer networks) 2. Computer network protocols.

3. Internetworking (Telecommunication) I. Title.

TH5105.543.S23 1999

004.6'2—dc21

99-36212

CIP

McGraw-Hill



A Division of The McGraw-Hill Companies

Copyright © 2000 by The McGraw-Hill Companies, Inc. All Rights Reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

1 2 3 4 5 6 7 8 9 0 DOC/DOC 9 0 4 3 2 1 0 9

ISBN 0-07-058098-7

The sponsoring editor for this book was Steven Elliot and the production supervisor was Clare Stanley. It was set in Century Schoolbook by D&G Limited, LLC.

Printed and bound by R. R. Donnelley & Sons Company.

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. For more information, please write to Director of Special Sales, McGraw-Hill, 11 West 19th Street, New York, NY 10011. Or contact your local bookstore.

Throughout this book, trademarked names are used. Rather than put a trademark symbol after every occurrence of a trademarked name, we use names in an editorial fashion only, and to the benefit of the trademark owner, with no intention of infringement of the trademark. Where such designations appear in this book, they have been printed with initial caps.

Information contained in this work has been obtained by The McGraw-Hill Companies, Inc. ("McGraw-Hill") from sources believed to be reliable. However, neither McGraw-Hill nor its authors guarantees the accuracy or completeness of any information published herein and neither McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that McGraw-Hill and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.



This book is printed on recycled, acid-free paper containing a minimum of 50 percent recycled de-inked fiber.

This book is dedicated to my wife, Nancy, for her consistent strength, direction, support, encouragement, understanding and for creating a life with me filled with love and affection.

ACKNOWLEDGMENTS

First and foremost I must thank my wife Nancy for her incredible support on this time consuming and lengthy project. She is the best. Secondly I must extend my appreciation to Steve Elliot, Executive Editor at McGraw-Hill, who throughout this project was able to keep me driven, focused, and excited about the handbook. I would also like to acknowledge Tony Costa CCIE# 4140 and Andrew Cassera CIT (CCIE in training) for their much needed contribution to this project. Their efforts and input were greatly appreciated. I would also like to thank Gary Stewart, Technical Illustrator with Cisco Systems Icons Document Resource Connection department for not getting totally annoyed with my request for hardware illustrations and Cisco Systems icons for this project. Not to be forgotten are the talented people who have made this Cisco Systems the definitive networking company. Their insight, dedication and fortitude have not only created a multibillion dollar company but have forged cultural and social change not seen since the availability of the automobile to the masses. All their efforts are greatly appreciated. Finally, I must thank all my children, Chelsea, Pamela, James, Meredith and Abigail for accepting my absence, irritability after many sleepless nights, and for the many times they have made me smile.

CONTENTS

| | | |
|------------------|---|-------------|
| | Foreword | xxix |
| Chapter 1 | Cisco IOS Software | 1 |
| | Benefits | 2 |
| | Scalability | 2 |
| | Adaptiveness | 3 |
| | Access Support | 3 |
| | Performance Optimization | 4 |
| | Management | 5 |
| | Security | 6 |
| | Packaging | 6 |
| | Feature Sets | 6 |
| | Feature Packs | 7 |
| | Features Supported | 7 |
| | Protocols | 7 |
| | Management | 12 |
| | Multimedia and QoS | 13 |
| | Secure Data Transmission | 14 |
| | Support for IBM Networking Environments | 14 |
| | IP Routing Protocols | 15 |
| | Bridging | 16 |
| | Packet Switching | 17 |
| | NetFlow Switching | 18 |
| | ATM | 18 |
| | Dial-on-demand Routing | 18 |
| | Access Server | 19 |
| | LAN Extension | 20 |
| Chapter 2 | Cisco Router Hardware | 23 |
| | Cisco Router Network Hierarchy | 24 |
| | Core | 24 |
| | Distribution | 26 |
| | Access | 26 |
| | Online Insertion and Removal (OIR) | 28 |
| | Supported Platforms | 28 |
| | OIR Process | 28 |
| | Exceptions to Using OIR | 28 |

| | |
|--|----|
| Cisco 12000 Series | 29 |
| Cisco 12004 Series | 30 |
| Cisco 12008 Series | 30 |
| Cisco 12012 Series | 32 |
| 12000 Usage | 32 |
| 12000 Switch Processors | 34 |
| 12000 Memory | 35 |
| 12000 Line Cards | 36 |
| 12000 Software Support | 42 |
| Cisco 7500 Series | 43 |
| Cisco 7505 Series | 44 |
| Cisco 7507 Series | 44 |
| Cisco 7513 Series | 45 |
| 7500 Usage | 46 |
| 7500 System Processors | 47 |
| 7500 Memory | 48 |
| 7200 Series | 49 |
| 7200 Usage | 51 |
| 7200 Network Processing Engine | 51 |
| 7200 Memory | 51 |
| 7000 Series | 52 |
| 7000 Usage | 53 |
| 7000 System Processors | 53 |
| 7000 Memory | 54 |
| Cisco 7x00 Series Interface Processors | 55 |
| 7x00 ATM Interface Processor (AIP) | 55 |
| 7x00 Channel Interface Processor 2 (CIP2) | 56 |
| 7x00 Channelized T3 Interface Processor (CT3IP) | 58 |
| 7x00 Ethernet Interface Processor (EIP) | 59 |
| 7x00 Fast Ethernet Interface Processor (FEIP) and FEIP2 | 59 |
| 7x00 FDDI Interface Processor (FIP) | 61 |
| 7x00 Fast Serial Interface Processor (FSIP) | 61 |
| 7x00 High-Speed Serial Interface(HSSI) Interface Processor (HIP) | 62 |
| 7x00 Multichannel Interface Processor (MIP) | 63 |
| 7x00 Packet OC-3 Interface Processor (POSIP) | 64 |
| 7x00 Service Provider MIP (SMIP) | 65 |
| 7x00 Standard Serial Interface Processor (SSIP) | 66 |
| 7x00 Token Ring Interface Processor (TRIP) | 67 |
| 7x00 Versatile Interface Processor 2 (VIP2) | 67 |

| | | |
|------------------|--|------------|
| | Cisco 7x00 Series Port and Service Adapters | 68 |
| | 7x00 ATM OC-3 | 69 |
| | 7x00 ATM-Circuit Emulation Services (ATM-CES) | 69 |
| | 7x00 100VG-AnyLAN | 70 |
| | 7x00 ISDN Basic Rate Interface (BRI) | 71 |
| | 7x00 Channelized T1/E1 ISDN PRI | 72 |
| | 7x00 Ethernet 10BaseT | 72 |
| | 7x00 Ethernet 10BaseFL | 73 |
| | 7x00 Fast Ethernet | 73 |
| | 7x00 Synchronous Serial | 74 |
| | 7x00 Single Port Molex 200-pin Receptacle | 75 |
| | 7x00 Synchronous Serial E1-G.703/G.704 | 75 |
| | 7x00 Token Ring | 75 |
| | 7x00 FDDI | 76 |
| | 7x00 HSSI | 77 |
| | 7x00 Compression Service Adapter | 77 |
| | 4000 Series | 78 |
| | 4000 Usage | 80 |
| | 4000 Processors | 80 |
| | 4000 Memory | 81 |
| | 3600 Series | 81 |
| | 3600 Usage | 90 |
| | 3600 Processors and Memory | 91 |
| | 2600 Series | 91 |
| | 2600 Usage | 93 |
| | 2600 Processor | 94 |
| | 2600 Memory | 94 |
| | 2500 Series | 94 |
| | 2500 Usage | 104 |
| | 2500 Processor and Memory | 104 |
| | 1600 Series | 105 |
| | 1600 Usage | 107 |
| | 1600 Processor and Memory | 107 |
| | 700M Family of Access Routers | 108 |
| Chapter 3 | Cisco Router Network Design | 113 |
| | The NetworkInfrastructure Lifecycle | 115 |
| | Executive Corporate Vision | 116 |
| | Gather Network Infrastructure Information | 116 |
| | Determine Current Network Requirements | 116 |

| | | |
|------------------|---|------------|
| | Assess Current Network Operational Processes | 116 |
| | Research Plans for New Applications | 116 |
| | Identify Networking Technologies | 117 |
| | Define a Flexible Strategic/Tactical Plan | 117 |
| | Developing an Implementation Plan | 117 |
| | Management and Review | 117 |
| | Design Criteria | 118 |
| | The Current LAN/Campus Network Trend | 118 |
| | WAN Design Trends | 121 |
| | Remote Network Trends | 122 |
| | Application Availability Versus Cost-Effectiveness | 123 |
| | Application Profile | 124 |
| | Cost-Efficiency | 126 |
| | Network Devices and Capabilities | 127 |
| | Bridging and Routing | 128 |
| | Switching | 129 |
| | Backbone Considerations | 131 |
| | Distributed Services | 133 |
| | Local Services | 135 |
| | Selecting Routing Protocol | 136 |
| | Network Topology | 137 |
| | Addressing and Route Summarization | 137 |
| | Route Selection | 138 |
| | The Concept of Convergence | 140 |
| | Network Scalability | 140 |
| | Security | 142 |
| Chapter 4 | IP Routing Protocol Design | 143 |
| | RIP, RIP2, and IGRP Network Design | 145 |
| | RIP, RIP2, and IGRP Topology | 146 |
| | RIP, RIP2, and IGRP Addressing and Summarization | 146 |
| | RIP, RIP2, and IGRP Route Selection and Convergence | 147 |
| | RIP, RIP2, and IGRP Network Scalability | 148 |
| | EIGRP Network Design | 148 |
| | EIGRP Topology | 149 |
| | EIGRP Addressing and Summarization | 149 |
| | EIGRP Route Selection | 149 |
| | EIGRP Convergence | 150 |
| | EIGRP Scalability | 151 |
| | EIGRP Scalability | 152 |
| | EIGRP Security | 152 |

| | | |
|------------------|---|------------|
| | OSPF Network Design | 152 |
| | OSPF Topology | 152 |
| | OSPF Addressing and Summarization | 154 |
| | OSPF Route Selection | 157 |
| | OSPF Convergence | 158 |
| | OSPF Scalability | 158 |
| | OSPF Security | 159 |
| Chapter 5 | Frame Relay Network Design | 161 |
| | Hierarchical Design of Frame Relay Internetworks | 162 |
| | Frame Relay Scalability | 164 |
| | Frame Relay Management | 164 |
| | Frame Relay Performance | 164 |
| | Frame Relay Network Topology | 165 |
| | Frame Relay Star Topology | 165 |
| | Frame Relay Fully Meshed Topology | 166 |
| | Frame Relay Partially Meshed Topology | 168 |
| | Frame Relay Fully Meshed Hierarchical Topology | 168 |
| | Frame Relay Hybrid Meshed Hierarchical Topology | 169 |
| | Broadcast Traffic Issues | 171 |
| | Performance Considerations | 171 |
| | Determining Maximum rate | 174 |
| | Committed Information Rate (CIR) | 174 |
| | FECN/BECN Congestion Protocol | 175 |
| | Virtual Subinterface and Multiprotocol Management | 175 |
| | SNA Support | 175 |
| | Boundary Network Node (BNN) | 176 |
| | Boundary Access Node (BAN) | 177 |
| | FRAS Host Support | 178 |
| Chapter 6 | ATM Internetworking Design | 181 |
| | LAN Emulation (LANE) | 182 |
| | LAN Emulation Client (LEC) | 183 |
| | LAN Emulation Configuration Server (LECS) | 184 |
| | LAN Emulation Server (LES) | 184 |
| | Broadcast Unknown Server (BUS) | 185 |
| | LANE Design Considerations | 185 |
| | Network Support | 186 |
| | Addressing | 186 |
| | LANE ATM Addresses | 186 |

| | | |
|------------------|--|------------|
| | Cisco's Method of Automatically Assigning ATM Addresses | 187 |
| | Using ATM Address Templates | 188 |
| | Rules for Assigning Components to Interfaces and Subinterfaces | 188 |
| | Redundancy in LANE Environments | 189 |
| | Data Exchange Interface (DXI) | 192 |
| | Supported Modes | 192 |
| | DXI Addressing | 193 |
| | Classical IP | 193 |
| | Multiprotocol over ATM (MPOA) | 194 |
| | Multiprotocol Client (MPC) | 195 |
| | Multiprotocol Server (MPS) | 196 |
| | MPOA Guidelines | 197 |
| | Bandwidth Support on Routers | 197 |
| | Configurable Traffic Parameters | 198 |
| Chapter 7 | Switched LAN Design | 199 |
| | Switched LAN Factors | 200 |
| | Broadcast Radiation | 201 |
| | Well-behaved VLANs | 202 |
| | Inter-VLAN Available Bandwidth | 202 |
| | Administrative Boundaries | 202 |
| | Cisco VLAN Implementation Support | 202 |
| | IEEE 802.10 | 205 |
| | IEEE 802.1d | 207 |
| | Inter-Switch Link (ISL) | 208 |
| | LAN Emulation | 209 |
| | Virtual Multi-homed Servers | 209 |
| | Switched LAN Topologies | 210 |
| | Scaled Switching | 211 |
| | Large Switching/Minimal Routing | 211 |
| | Distributed Routing/Switching | 213 |
| Chapter 8 | SRB/RSRB Network Design | 215 |
| | Steps to Effective SRB Design | 216 |
| | Determine SRB Protocols | 216 |
| | Determine Parallel Link Requirements | 217 |
| | Specify an Appropriate RSRB Encapsulation Technique | 217 |
| | Determine WAN Frame Size Requirement | 218 |

| | |
|--|-----|
| Determine Local Acknowledgment Requirement | 218 |
| Select the Appropriate IP Routing Protocol | 218 |
| Explorer Packet Control | 218 |
| NetBIOS Traffic Management | 219 |
| Typical SRB Topologies | 219 |
| Hierarchical | 220 |
| Distributed | 220 |
| Flat | 220 |
| Virtual Ring Concept | 224 |
| Multiport Bridging | 225 |
| Redundant Star Topology | 225 |
| Fully Meshed Topology | 227 |
| Hierarchical Virtual Ring Topology | 227 |
| Explorer Packet on Virtual Rings | 227 |
| Proxy Explorer | 229 |
| NetBIOS Broadcast Control | 230 |
| NetBIOS Name Caching | 231 |
| NetBIOS Datagram Broadcast | 232 |
| NetBIOS Broadcast Throttling | 232 |
| NetBIOS Broadcast Dampening | 233 |
| Remote SRB Encapsulation Techniques | 233 |
| Direct | 234 |
| TCP Encapsulation | 235 |
| Fast Sequenced Transport (FST) Encapsulation | 236 |
| Parallel WAN Link Issues | 236 |
| Process Switching | 238 |
| Fast Switching | 238 |
| Effect of IP Routing Protocols | 238 |
| Local Acknowledgment | 239 |
| Design Recommendations | 240 |
| IP Routing Protocols and SRB | 240 |
| Link Failure Effects | 241 |
| IGRP Network Design | 241 |
| Enhanced IGRP Network Design | 242 |
| OSPF Network Design | 242 |
| Queuing and Prioritization | 242 |
| Priority Queuing | 243 |
| Custom Queuing | 243 |
| SAP Prioritization | 243 |
| LU Address Prioritization | 244 |
| SAP Filters for WAN Links | 244 |

| | | |
|-------------------|---|------------|
| Chapter 9 | DLSw+ Network Design | 245 |
| | DLSw Standard | 246 |
| | Peer Connections | 248 |
| | Exchange of Capabilities | 248 |
| | Circuit Establishment | 249 |
| | Flow Control | 249 |
| | Cisco DLSw+ | 250 |
| | Peer Group Concept | 252 |
| | Explorer Firewalls | 254 |
| | Fault Tolerance | 254 |
| | Load Balancing | 255 |
| | Cisco DLSw+ Transport | 256 |
| | TCP Encapsulation | 256 |
| | Fast Sequenced Transport (FST) | 258 |
| | Direct | 258 |
| | DLSw Lite | 259 |
| | Encapsulation Overhead | 259 |
| | Cisco's Enhanced Modes of Operation | 259 |
| | Dual Mode | 260 |
| | Standard Mode | 260 |
| | Enhanced Mode | 260 |
| | Availability Configurations | 261 |
| | Load Balancing | 261 |
| | Redundancy | 263 |
| | Backup Peers | 263 |
| | Performance Features | 265 |
| | Port Lists | 265 |
| | Peer Groups with Border Peers and On-demand Peers | 265 |
| | Dynamic Peers | 267 |
| Chapter 10 | APPN Network Design | 269 |
| | Cisco Support of APPN | 272 |
| | Deciding Factors on Using APPN | 273 |
| | The Dominant Protocol Is SNA | 274 |
| | Class of Service (COS) | 274 |
| | Office-to-Office Connectivity | 275 |
| | Non-APPN-capable SNA Resources | 277 |
| | Dependent Logical Unit Requester/Server | 277 |
| | Network Node Placement | 277 |

| | |
|--|------------|
| Performance Considerations | 278 |
| Topology Database Updates | 279 |
| Connection Network | 280 |
| Number of CP-CP Sessions | 281 |
| Number of Network Nodes | 284 |
| LOCATE Search | 285 |
| Recovery Techniques | 286 |
| Secondary WAN Link | 286 |
| Dual Routers/Dual WAN Links | 286 |
| High-Performance Routing | 287 |
| SSCP Takeover of DLUR | 288 |
| Queuing and Prioritization | 289 |
| APPN Buffer and Memory Management | 289 |
| | |
| Chapter 11 ISDN and Dial-On-Demand Routing (DDR) Design | 293 |
| Site Options | 294 |
| Central Office Switch Considerations | 295 |
| PRI and BRI | 297 |
| DDR Model | 298 |
| DDR Dialer Cloud | 299 |
| IP Addressing | 300 |
| Topology | 300 |
| Point-to-Point | 300 |
| Hub and Spoke | 301 |
| Fully Meshed | 302 |
| Dial Service Considerations | 302 |
| Data Encapsulation | 302 |
| Synchronous Lines | 303 |
| Asynchronous | 304 |
| ISDN | 304 |
| Dialer Rotary Groups | 304 |
| Dialer Profiles | 305 |
| Dialer Maps | 305 |
| Routing Packets | 305 |
| Static Routing | 306 |
| Dynamic Routing | 306 |
| Using Passive Interfaces | 306 |
| Split Horizons | 307 |
| Dynamically Connected Routes | 307 |
| Snapshot Routing | 307 |
| DDR as Dial Backup | 309 |

| | |
|---|------------|
| Connection Triggering | 309 |
| Bandwidth | 310 |
| Interesting Packets | 310 |
| Controlling Routing Updates as Triggers | 311 |
| Access Lists | 312 |
| Security | 313 |
| Callback | 313 |
| Screening | 313 |
| Chapter 12 Preparing the Cisco Router | 315 |
| Determining the Proper IOS Code | 316 |
| Locate IOS Using Cisco Connection Online (CCO) | 319 |
| CCO Software Center | 320 |
| Downloading IOS to the TFTP Server | 321 |
| Loading IOS on the Router | 321 |
| Enter Privileged Mode on the Terminal Interface | 322 |
| Issue Copy to/from a TFTP Server | 323 |
| Issue Copy to/from the FTP Server | 326 |
| Loading CIP or CPA Microcode on a Cisco 7000/7200/ 7500 Router | 328 |
| Locating CIP/CPA Microcode on the CCO Software Center | 329 |
| Enter config Parameters to Load Microcode | 330 |
| Router Basic Configuration and IOS Commands | 333 |
| Setting up EXEC and Privileged Mode Access | 333 |
| BOOT SYSTEM List | 336 |
| BOOT CONFIG List | 337 |
| Assigning a Name to the Router | 339 |
| Enabling a DNS Search and Assigning a Domain | 339 |
| Specifying SNMP for Router Management | 340 |
| Using the Banner Command | 347 |
| Chapter 13 IP Configuration | 349 |
| Defining Subnets on the Router Interfaces | 350 |
| When to Assign Multiple IP Addresses to an Interface | 354 |
| Maximizing Network Address Space Using Subnet Zero | 356 |
| Routing Same Network Packets Without a Routing Table Entry | 357 |

| | |
|---|-----|
| Enable IP on a Serial Interface Without Specifying an IP Address | 357 |
| IP Address Mapping | 359 |
| Define a Static ARP Cache | 362 |
| Supporting Multiple ARP Encapsulation Types on an Interface | 363 |
| Disable Proxy ARP | 363 |
| Enabling the Use of IP Host Names in Cisco IOS | |
| Commands | 363 |
| Static Definition of IP Host Names and IP Addresses | 364 |
| Enable the Use of DNS Services | 364 |
| Disable IP Routing | 366 |
| Routing Assistance When IP Routing Is Disabled | 366 |
| Bridging IP, Instead of Routing IP | 369 |
| Controlling and Managing Broadcast Packets | 370 |
| Enable Directed Broadcast-to-Physical Broadcast Translation | 371 |
| Configure IP Services | 373 |
| Disable ICMP Unreachable Messages | 373 |
| Disable ICMP Redirect Messages | 374 |
| Disable ICMP Mask Reply Messages | 374 |
| Support for Fragmenting Large IP Packets | 375 |
| Filter IP Packets Using Access Lists | 376 |
| Create Standard Access Lists Using Numbers and Names | 378 |
| Extended Access Lists | 381 |
| Fault-Tolerant Routing of IP Packets | 390 |
| IP Performance Tuning | 396 |
| Compressing TCP Headers | 396 |
| Enable TCP Path MTU Discovery | 397 |
| Performance Enhancements Through Switching | |
| Features | 398 |
| Enabling Process Switching | 398 |
| Enabling Fast Switching | 399 |
| Enabling Fast Switching on the Same Interface | 399 |
| Enabling Autonomous Switching | 399 |
| Enabling Silicon Switch Engine (SSE) Switching | 400 |
| Optimum Fast Switching | 400 |
| NetFlow Switching | 400 |
| Distributed Switching | 403 |

| | | |
|-------------------|--|------------|
| Chapter 14 | Defining RIP Routing Protocol | 405 |
| | The Basics of RIP | 406 |
| | Why Use RIP as the Routing Protocol? | 407 |
| | Defining RIP as a Routing Protocol on the Router | 408 |
| | Allow Point-to-Point (NBMA) Updates for RIP | 411 |
| | Specifying the Version of RIP | 416 |
| | Enabling RIP Version 2 Authentication | 419 |
| | Disable RIP Version 2 Route Summarization | 423 |
| | Disabling the Validation of Source IP Addresses | 426 |
| | Reducing Routing Loops with Split-horizon | 428 |
| | Tuning RIP Update Packet Delays | 428 |
| | RIP Updates and the Affect on Bandwidth | 429 |
| | | |
| Chapter 15 | Configuring IGRP Routing Protocol | 431 |
| | Managing IGRP Updates and Route Advertisements | 433 |
| | The Holddown Algorithm for IGRP | 434 |
| | The Split-horizon Algorithm for IGRP | 434 |
| | The Poison-reverse Updates Algorithm for IGRP | 437 |
| | Defining IGRP as a Routing Process | 437 |
| | Using Unicast IGRP Routing Updates | 441 |
| | Increasing Throughput and Reliability Using Unequal- Cost Paths | 445 |
| | Altering IGRP Routing and Metric Computations | 447 |
| | Decreasing IGRP Route Convergence | 451 |
| | Tuning IGRP Route Convergence | 453 |
| | Controlling the Logical Size of an IGRP Network | 454 |
| | Validate Source IP Addresses | 455 |
| | IGRP Updates and the Affect on Bandwidth | 456 |
| | | |
| Chapter 16 | Configuring Enhanced IGRP (EIGRP) Routing Protocol | 459 |
| | Enable EIGRP as a Routing Protocol | 461 |
| | Migrating to EIGRP from IGRP | 465 |
| | Monitoring Neighbor Adjacency Changes | 467 |
| | Managing EIGRP Bandwidth Utilization | 469 |
| | Modifying EIGRP Metric Weights | 472 |
| | Routing Between Disconnected Networks with EIGRP | 476 |
| | Summarizing Routes for Advertisements out Specific Interfaces | 478 |

| | | |
|-------------------|--|------------|
| | Tuning Hello Packet and the Hold Time intervals | 479 |
| | Split-Horizon and EIGRP | 481 |
| | MD5 Authentication with EIGRP | 482 |
| Chapter 17 | Configuring OSPF Routing Protocol | 487 |
| | OSPF with Cisco IOS | 490 |
| | Specifying OSPF on Cisco Routers | 491 |
| | Creating the OSPF Routing Process | 492 |
| | DR and BDR Election Using a Loopback Interface | 497 |
| | Multi-Area OSPF Networks | 500 |
| | OSPF Route Summarization | 501 |
| | Stub, Totally Stubby, and Not-So-Stubby-Area (NSSA) OSPF Areas | 507 |
| | Configuring a Stub and Totally Stubby Areas | 509 |
| | Configuring a NSSA | 511 |
| | Using OSPF Virtual Links | 514 |
| | Configuring Virtual Links | 514 |
| | Non-Broadcast and Broadcast Network Configurations | 517 |
| | Configuring OSPF Point-to-Multipoint | 518 |
| Chapter 18 | Configuring BGP Routing Protocol | 525 |
| | Exterior and Interior BGP Sessions | 527 |
| | Path Selection Under BGP | 529 |
| | Defining the BGP Process | 530 |
| | Enabling the BGP Process | 530 |
| | Establishing BGP Peer Connections | 531 |
| | Forcing a Loopback Interface as the BGP Neighbor | 534 |
| | Multihop EBGP Connections | 539 |
| | Controlling Routing Information Using Route Maps | 542 |
| | Redistributing IGP Routes into BGP | 546 |
| | Nexthop Attribute and Multiaccess/NBMA Networks | 548 |
| | Forcing BGP to Prefer a IGP Route Using the Backdoor Command | 552 |
| | Synchronization of Routes | 553 |
| | Best Path Selection Using the Weight Attribute | 555 |
| | Forcing a Preferred AS Exit Path | 556 |
| | Influencing the Preferred Path into an AS | 558 |
| | Grouping Destinations That Share a Common Attribute | 560 |
| | Route and Path Information Filtering | 564 |

| | | |
|-------------------|--|------------|
| | Manipulating the Path Information Order | 571 |
| | Using BGP Peer Groups | 572 |
| | Aggregate Addresses | 574 |
| | Grouping Multiple Autonomous Systems into a BGP Confederation | 579 |
| | Reducing IBGP Peering Using Route Reflectors (RR) | 582 |
| | Managing Unstable Routes | 589 |
| Chapter 19 | Route Redistribution | 593 |
| | Understanding Route Redistribution | 594 |
| | The Redistribution Router Configuration Command | 595 |
| | Selecting the Best Path Based on a Routing Protocol | 598 |
| | Altering Trusted or Believable Routes Using the Distance Command | 598 |
| | Modifying the Seed Metric Using the default-metric Command | 600 |
| | Filtering Redistributed Routes Using the distribute- list Command | 603 |
| | Redistribution Considerations | 605 |
| | Redistribution Examples | 606 |
| | RIP V1-to-OSPF Redistribution | 606 |
| | IGRP-to-EIGRP Redistribution | 612 |
| | RIP-to-EIGRP Redistribution | 616 |
| Chapter 20 | Defining ATM (LANE, Classical IP, and MPOA) | 619 |
| | Configuring LANE | 620 |
| | Initial Configuration for LANE | 620 |
| | Configuring a LAN Emulation Client (LEC) | 622 |
| | Configuring a LAN Emulation Configuration Server (LECS) | 629 |
| | Configuring a LAN Emulation Server (LES) | 635 |
| | Configure ATM over a Serial Interface | 638 |
| | Configuring Classical IP | 640 |
| | Configuring Multiprotocol over ATM (MPOA) | 644 |
| | Configuring a MPOA Client (MPC) | 646 |
| | Configuring a MPOA Server (MPS) | 649 |
| Chapter 21 | Defining Frame Relay | 655 |
| | A Simple Frame Relay Configuration | 656 |
| | Dynamic and Static Addressing | 659 |

| | |
|---|------------|
| Frame Relay Subinterfaces | 664 |
| Point-to-Point Frame Relay Subinterfaces | 664 |
| Multipoint Frame Relay Subinterfaces | 664 |
| Specifying a Frame Relay Subinterface | 666 |
| Hub and Spoke Configurations | 670 |
| Dynamic Addressing for IP-Only Connectivity | 670 |
| Static Addressing over a Multipoint Configuration | 672 |
| Traffic Shaping on Frame Relay | 674 |
| Transparent Bridging | 678 |
| Managing Performance Problems Using the Broadcast Queue | 681 |
| Chapter 22 Internet-working Legacy Systems | 683 |
| Serial Tunneling (STUN) | 685 |
| Source-Route Bridging Configuration | 695 |
| Data Link Switching Configuration Plus (DLSw+) | 704 |
| SDLC to LLC2 (SDLLC) | 722 |
| Frame Relay BNN and BAN Connectivity | 731 |
| APPN Support on Cisco Routers | 741 |
| SNA Support Using a Channel-Attached Cisco Router | 750 |
| IP Connectivity Using a Channel-Attached Cisco Router | 761 |
| Chapter 23 Defining Novell Networks | 765 |
| IPX Processing | 766 |
| Configuring for IPX Routing | 766 |
| Routing IPX | 769 |
| IPX RIP | 769 |
| NLSP Routing | 772 |
| IPX EIGRP Routing | 782 |
| IPX Static Routes | 791 |
| IPX Route Redistribution | 792 |
| Bridging IPX | 793 |
| Transparent Bridging | 794 |
| Encapsulated Bridging | 794 |
| Chapter 24 Configuring ISDN | 797 |
| Configuring BRI Service | 798 |
| Configuring Primary Rate Interface (PRI) Service | 802 |
| Dial-On-Demand routing (DDR) with ISDN | 805 |
| Using Caller ID | 809 |

Contents

| | |
|----------------------------|------------|
| ISDN Callback | 810 |
| SNAPSHOT Routing with ISDN | 812 |
| Using ISDN for dial backup | 817 |
| Appendix A | 821 |
| Appendix B | 827 |
| Appendix C | 831 |
| Appendix D | 833 |
| Glossary | 837 |
| Index | 873 |

About the Author

George C. Sackett is Managing Director of NetworX Corporation, a New Jersey professional services company which develops corporate data networks for the telecommunications, entertainment, medical, financial, and transportation industries. The author and coauthor of other McGraw-Hill titles on networks, he has more than 15 years of technical and managerial experience with corporate data networks.

About the Reviewers

As the leading publisher of technical books for more than 100 years, McGraw-Hill prides itself on bringing you the most authoritative and up-to-date information available. To ensure that our books meet the highest standards of accuracy, we have asked a number of top professionals and technical experts to review the accuracy of the material you are about to read.

We take great pleasure in thanking the following technical reviewers for their insights:

Jim Bender is currently with RPM Consulting, an internetworking consulting company. He has twenty-two years experience in IS/IT and has held positions as a mainframe Cobol/Assembler programmer, a systems programmer supporting the operating systems and network, and Manager of Technical Support. He has worked with Cisco routers supporting IBM, Novell, and Microsoft networks.

E. Gary Hauser, CCIE (#4489), CCSI(#99166), CCDA, Master CNE, CNE 4.x, 3.x novell is a Cisco certified systems instructor with Chesapeake Computer Consultants, Inc. based in San Jose, California. He currently teaches the BCRAN, ACRC, ICRC, and ICND courses nationwide as well as internal training for the Cisco SE technical staff. Recently he has been doing course development with Cisco for the new curriculum. He has over 11 years' experience in the industry, holding various positions at the Systems Engineer and Network manager levels for ADC telecommunications, SURAnet, and a large Washington, D.C.-based health care provider, prior to becoming an Instructor/Consultant for the Annapolis, MD-based Chesapeake Computer Consultants, Inc.

Foreword

The phenomenal growth of TCP/IP and the Internet has its foundation in Cisco Systems routers. Cisco Systems has anywhere from 60–80% share of the router market. The handbook uses real world network configurations. Each configuration is defined, analyzed and described so that any networking professional will be able to configure a Cisco router in their network. The handbook explores the typical examples for router networks and explains in detail the parameters most commonly used in configuring a Cisco router.

The intended readership for this book covers a wide array of networking professionals. These include traditional networking professionals, new to the world of routing and experienced routing professionals needing a quick definitive source for configuring routers.

Cisco routers provide not only TCP/IP network connectivity but have become a means for transporting many network protocols. This ability has made the configuration of Cisco routers convoluted and complex at times. The handbook guides network professionals through these complex configurations.

The handbook is organized into two sections. The first section covers network design issues for the majority of networking topologies found in corporate networks. The second half of the handbook discusses in detail the following topics:

- Overview of Cisco Internetwork Operating System (IOS).
- Connectivity to multiple network infrastructures: Ethernet, Token-ring, T1, ATM, Frame-Relay, ISDN.
- Configuring networks with routing protocols: RIP, RIP-2, IGRP, E-IGRP, OSPF, BGP and Novell IPX.
- Discussions on incorporating routers into ATM, VLANs, ELANs and support for Multiprotocol over ATM (MPOA).
- Internetworking legacy mainframe systems using RSRB, DLsw+, TCP/IP, Channel Interface Processors (CIP), and IBM APPN.




The handbook may not cover every configuration command or address all the possible network topologies in use. It is the intention of this handbook to become a reference for the majority of network topologies that every networking professional will encounter at some point in their career.

I sincerely hope the handbook meets your expectation. You can send comments to me through e-mail using gsackett@networxcorp.com. I look forward to reading them.

CHAPTER

1

Cisco IOS Software



We have all heard the saying, “It’s what’s inside that counts.” In the world of networking, Cisco’s Internetwork Operating Systems (IOS) has taken that saying to heart. The very core of Cisco Systems’ phenomenal success is the breadth of services provided by the Cisco IOS software.

No two networks are exactly alike. Connectivity requirements differ between healthcare and manufacturing, entertainment and shipping, finance and telecommunications—each of which has different security issues. Each requires the capability to scale with reliability and manageability. The Cisco IOS software has proven to meet these criteria and to build on new requirements due to its flexibility in meeting the rapidly changing network requirements of all businesses.

Benefits

Cisco IOS software provides a foundation for meeting all the current and future networking requirements found in today’s complex services-driven business environments. Businesses rely heavily on generating income from their network infrastructure. Cisco IOS software has the broadest set of networking features primarily based on international standards allowing Cisco products to interoperate with disparate media and devices across an enterprise network. Most importantly, Cisco IOS software enables corporations to deliver mission-critical applications seamlessly between various computing and networking systems.

Scalability

The network infrastructure for every corporation must be flexible to meet all the current and future internetworking requirements. Cisco IOS software uses some proprietary but also adheres to international standards for congestion avoidance using scalable routing protocols. These routing protocols allow a network using Cisco IOS to overcome network protocol limitations and deficiencies inherent in the protocols’ architectures. Additional features in scaling an efficient use of bandwidth and resources include the ability of the IOS software in detailed packet filtering for reducing “chatty” protocol traffic as well as reducing network broadcasts through timers and helper addresses. All these features and more are available with the goal of reducing network traffic overhead, thereby maintaining an efficient yet effective network infrastructure.

Adaptiveness

Network outages occur frequently in corporate networks. However, these outages do not often affect the flow of business due to the reliability and adaptiveness of the policy-based IOS software routing features. Using routing protocols, each Cisco router can dynamically decide on the best route for delivering packets through the network around outages, thereby providing reliable delivery of information. The prioritization of packets and services enables Cisco routers to adapt to bandwidth constraints due to outages or high bandwidth utilization. IOS software load-balances traffic throughput over various network connections, preserving bandwidth and maintaining network performance.

The concept of virtual local area networks (LANs) has become a reality for many corporate networks. Cisco routers have the capability to participate in these virtual LANs using emulated LAN functions for physical LAN extensions and ATM LAN Emulation (LANE) services. These are just two of the many newer networking technologies incorporated into the IOS software feature set, enabling networks to implement newer technologies without the added expense of new hardware.

Access Support

The Cisco IOS software access support encompasses remote access and protocol translation services. These services provide connectivity to

- Terminals
- Modems
- Computers
- Printers
- Workstations

Various network configurations exist for connecting these network resources over LANs and wide area networks (WANs). LAN terminal service support is as follows:

- TCP/IP support for Telnet and rlogin connections to IP hosts.
- TN3270 connections to IBM hosts.
- LAT connections to DEC hosts.

For WANs, Cisco IOS software supports four flavors of server operations:

- Connectivity over a dial-up connection supporting
 - AppleTalk Remote Access (ARA)
 - Serial Line Internet Protocol (SLIP)
 - Compressed SLIP (CSLIP)
 - Point-to-Point Protocol (PPP)
 - Xremote, Network Computing Device's (NCD) X Window System terminal protocol
- Asynchronous terminal connectivity to a LAN or WAN using network and terminal emulation software supporting Telnet, rlogin, DEC's Local Area Transport (LAT) protocol, and IBM TN3270 terminal protocol.
- Conversion of a virtual terminal protocol into another protocol, such as LAT-TCP or TCP-LAT communication between a terminal and a host computer over the network.
- Support for full Internet Protocol (IP), Novell Internet Packet Exchange (IPX), and AppleTalk routing over dial-up asynchronous connections.

Performance Optimization

Optimizing networks requires network equipment to dynamically make decisions on routing packets in a cost-effective manner over the network. Cisco IOS software has two features that can greatly enhance bandwidth management, recovery, and routing in the network. These two features are dial-on-demand access (DDA) and dial-on-demand routing (DDR).

DDA is useful in several scenarios. These are

- Dial backup
- Dynamic bandwidth

In many instances, connectivity to a location fails because of a modem, DSU/CSU failure, or the main telecommunications line to the office is disrupted in some way. A good network design has a backup solution for this type of outage. Using DDA, a router can sense the line outage and perform a dial backup connection over a switched serial, ISDN, T1, or frame relay. In this manner, the office maintains connectivity to the WAN with minimal

downtime. The DDA function monitors the primary line for activation and can cut back to the primary connection automatically if so desired.

DDA features the capability to determine a low and high bandwidth watermark on the permanent lines. This feature allows the addition of a temporary bandwidth to another location to meet throughput and performance criteria. The IOS monitors the permanent line for high bandwidth utilization. If the bandwidth reaches the defined threshold, DDA is enabled to add extra bandwidth to the remote location of the permanent line. IOS continues to monitor the bandwidth for utilization to fall under the threshold for a period of time. Once the low watermark is reached, IOS disconnects the DDA line. Using DDA in this fashion enables the IOS to maintain performance criteria between the two locations.

DDR allows Cisco routers to create temporary WAN connections based on interesting packets. IP, Novell IPX, X.25, Frame Relay, and SMDS destination addresses can be specified under DDR as interesting packets. Once the router interprets the packet and determines it is an interesting packet, it performs the dialup connection to the destination network specified in the packet that corresponds to the DDR configuration. In this way, connectivity to remote locations is provided on a temporary basis, thereby saving network connectivity costs.

Management

Cisco IOS software supports the following protocols:

- The two versions of Simple Network Management Protocol (SNMP) for IP-based network management systems
- The Common Management Interface Protocol (CMIP)/Common Management Interface Service (CMIS) for OSI-based network management systems
- IBM Network Management Vector Transport (NMVT) for SNA-based network management systems

These management protocols are pertinent to the type of network supported by the Cisco router. The IOS itself has the capability to perform configuration management services as well as monitoring and diagnostics services using the IOS command interface.

Cisco Systems has a suite of network management tools under the name of CiscoWorks. CiscoWorks works with Cisco IOS for change, configuration, accounting, performance, and fault management disciplines.

Security

Cisco IOS software supports many different types of security capabilities. Some of these, such as filtering, are not usually thought of as a security feature. Filtering, for example, was actually the first means of creating the now infamous firewall techniques for corporate connectivity on the Internet, prior to actual commercial offerings.

Secondly, filtering can be used to partition networks and prohibit access to high-security server networks. The IOS has the capability to encrypt passwords, authenticate dial-in access, require permissions on changing configurations, and provide accounting and logging to identify unauthorized access.

The IOS supports standard authentication packages for access to the router. These are RADIUS and TACACS+. Each security package requires unique user identification for access to the router. These security packages offer multilevel access to IOS command interface functions.

Packaging

The ordering of Cisco IOS software has been streamlined into feature sets. Prior to IOS Version 11.2, the IOS software was built based on the router requirements. A second enhancement to the delivery of IOS software is the use of feature packs. Feature packs allow you to order the IOS software images and a Windows 95 utility to load the image on the router.

Feature Sets

Each feature set contains a standard offering. However, options are provided to enable the IOS software to meet more specific needs. Each hardware platform has a feature set. For the most part, all the routers share the same feature sets. The sets are broken down into three categories:

- Basic: The basic feature set for the platform.
- Plus: The basic feature set plus added features depending on the platform.
- Encryption: 40-bit (Plus 40) or 56-bit (Plus 56) data encryption feature sets with the basic or plus feature set.

The list of features and feature sets and the platforms supporting them are found in Appendix A.

Feature Packs

IOS Release 11.2 introduced software feature packs, which offer a means for receiving all materials including software images, loading utilities, and manuals on CD-ROMs. Each feature pack contains two CD-ROMs. The software CD-ROM contains

- IOS software images
- AS5200 modem software images
- Windows 95 software installer program

A second CD-ROM is included, providing the Cisco IOS software documentation reference library. The remaining documentation provided by the feature pack includes an instruction manual for using the Windows 95 software installer program, release notes for the IOS release included on the software CD-ROM, and the software license.

Features Supported

All the features found in the matrices of Appendix A are applicable to each router and access server platform. These features cover a wide range of services and functions to take into account old, current, and future network configurations.

Protocols

Cisco IOS supports a wide array of networking protocols. Of these protocols, Transmission Control Protocol/Internet Protocol (TCP/IP) is by far the most widely used.

TCP/IP Cisco IOS software supports the following TCP/IP features:

- IP access lists
- IP Security Option (IPSO)

- IP accounting
- Simple Network Management Protocol (SNMP)
- Serial Line Interface Protocol (SLIP)
- Address Resolution Protocol (ARP)
- Reverse Address Resolution Protocol (RARP)
- Domain Name System (DNS) support
- Internet Common Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- User Datagram Protocol (UDP)
- Telnet
- TN3270
- Trivial File Transfer Protocol (TFTP)

Release 10 and 10.3 of IOS introduced new features to already existing standards that have given Cisco routers the capability to provide higher levels of security, greater availability, and increased network scalability. Among these features are

- Hot Standby Router Protocol (HSRP) and Multigroup HSRP
- Next Hop Resolution Protocol (NHRP)
- Department of Defense Intelligence Information System Network Security for Information Exchange (DNSIX) extended IPSO
- Type of Service (TOS) queuing
- Cisco Discovery Protocol (CDP)
- Border Gateway Protocol (BGP) Communities

With the introduction of Release 11 and 11.1, the Cisco IOS software enhances router functionality in the areas of security, performance, and routing services. The major enhancements for these releases are

- Route Authentication with Message Digest 5 (MD5) encryption algorithm
- IP Access Control List (ACL) Violation Logging
- Policy-based routing
- Weighted fair queuing
- NHRP on IPX
- Fast Install for Static Routers

- Fast Switched GRE
- RIPV2

Release 11.2 implements more routing protocol enhancements, IP address translation features, and access control list usability. The major features introduced are

- On-Demand Routing (ODR) for stub routers
- OSPF On-Demand Circuit (RFC1793)
- OSPF Not-So-Stubby-Area (NSSA)
- BGP4 enhancements
 - Soft Configuration
 - Multipath
 - Prefix filtering with inbound route maps
- Network Address Translation (NAT)
- Named IP access control list
- Integrated routing and bridging (IRB)

ISO CLNS The Open Systems Interconnection (OSI) reference model implements the International Organization for Standardization's (ISO) Connectionless Network Service (CLNS) as the network layer protocol. Cisco IOS fully supports the forwarding and routing of ISO CLNS. The ISO standards and Cisco implemented features supported by Cisco IOS are

- ISO 9542 End System-to-Intermediate System (ESIS) routing protocol
- ISO 8473 Connectionless Network Protocol (CLNP)
- ISO 8348/Ad2 Network Service Access Points (NSAP)
- ISO 10589 Intermediate System-to-Intermediate System (IS-IS) routing protocol
- DDR for OSI/CLNS
- Connection-Mode Network Service (CMNS) for X.25 using NSAP

DECnet Phase IV and Phase V Cisco routers have supported DECnet for some time. IOS software has full functional support of local and wide area DECnet Phase IV and Phase V routing on all media types. Currently, Cisco IOS supports these enhanced DECnet features:

- DECnet dial-on-demand (DDR)
- Dynamic DECnet Route Advertisements
- DECnet Host Name-to-Address Mapping
- Target Address Resolution Protocol (TARP) support over SONET

Novell IPX Since IOS Release 10.0, Cisco IOS provides complete IPX support. Beginning with Release 10.3, IOS enhancements for Novell have centered on performance, management, security, and usability. These enhancements are

- Novell Link State Protocol (NLSP)
- IPXWAN 2.0
- IPX floating static routes
- SPX spoofing
- Enhanced IGRP to NLSP route redistribution
- Input access lists
- Per-host load-balancing
- NLSP route aggregation
- Raw FDDI IPX encapsulation
- IPX header compression
- Display SAP by name
- IPX ACL violation logging
- Plain English IPX access lists

AppleTalk Phase 1 and Phase 2 AppleTalk has been a long-standing supported protocol on Cisco IOS software. Extended and non-extended networks under AppleTalk Phase 2 are supported. Cisco IOS routes AppleTalk packets over all media types. The AppleTalk features implemented by Cisco IOS are

- MacIP
- IPTalk
- SNMP over AppleTalk
- Routing Table Maintenance Protocol (RTMP)
- AppleTalk Update-Based Routing Protocol (AURP)
- AppleTalk over enhanced IGRP

- Inter-enterprise routing
- AppleTalk Name Binding Protocol (NBP) filtering
- AppleTalk floating static routes
- Simple multicast routing protocol (SMRP)
- AppleTalk load-balancing
- SMRP fast switching

Banyan VINES Banyan's Virtual Integrated Network Service (VINES) is supported on all media types with Cisco IOS software. The VINES routing protocol itself automatically determines a metric for delivering routing updates. This metric is based on the delay set for the interface. Cisco IOS enhances this metric by allowing you to customize the value for the metric. Other enhancements and features supported on Banyan VINES using Cisco IOS are

- Address resolution in response to address requests and broadcast propagation
- MAC-level echo support to Ethernet, IEEE 802.2, Token Ring, and FDDI
- Name to address mapping for VINES host names
- Access list filtering of packets to or from specific networks
- Routing Table Protocol (RTP)
- Sequenced Routing Update Protocol (SRTP)
- VINES DDR
- Floating static routes

Xerox Network System (XNS) XNS is the foundation for the Novell IPX protocol. As such, Cisco IOS supports an XNS-routing protocol subset of the XNS protocol stack. XNS is supported on

- Ethernet
- FDDI
- Token Ring
- Point-to-point serial lines using HDLC
- Link Access Procedure Balanced (LAPB)
- X.25 Frame relay
- SMDS networks

Apollo Domain Apollo workstations use the Apollo Domain routing protocol. Cisco IOS supports packet forward and routing of this protocol on

- Ethernet
- FDDI
- HDLC
- X.25 encapsulation

HP Probe HP Probe is a protocol used by HP devices that provide machine name resolution to the physical IEEE 802.3 address. Cisco routers acting as HP Probe Proxy servers on IEEE 802.3 LANs allow the router to resolve the machine name to the IEEE 802.3 address, eliminating the need for a separate server on each IEEE 802.3 LAN and saving corporate resources.

Multiring Cisco IOS supports the framing of Layer 3 protocol packets in Source Route Bridging packets using the Multiring protocol. Multiring is primarily used for Token Ring networks.

Management

Cisco IOS software supports the three network management schemas:

- SNMP
- CMIP/CMIS
- IBM NMVT

These network management schemas used by network management applications execute on workstations, minicomputers, or mainframes. For the most part, they use a client/server type of architecture between the router and the management system.

IOS Release 11.2 introduced the capability to manage Cisco routers using HyperText Transfer Protocol (HTTP) from Web browsers. HTTP utilizes HyperText Markup Language (HTML) for navigating Web pages from a browser. All Cisco routers in Release 11.2 or higher have the capability of presenting a home page to a Web browser. The default home page allows you to use IOS command line interface commands through Web-like hot links. This home page is modifiable to meet the needs of any router or organization.

Web-based interfaces have been available for the Cisco 700, 1000, and 1600 access routers since IOS Release 11.0. This Web-based application on the Cisco access product line is called ClickStart. The ClickStart interface presents at installation an initial setup form, guiding the operator through router configuration. Once the router is configured and connected to the network, it is manageable from any central location.

Multimedia and QoS

The advent of higher bandwidth and technologies enabling the integration of audio, video, and data on the same network medium have given rise to the need for supporting multimedia applications with guaranteed service. Cisco IOS Release 11.2 meets the quality of service (QoS) requirement of the following multimedia applications:

- Resource Reservation Protocol (RSVP)
- Random Early Detection (RED)
- Generic Traffic Shaping.

RSVP is an IETF standard that enables applications to dynamically reserve network resources (bandwidth, in other words) from end-to-end. Video or audio feeds over the network can now coexist with data traffic without the need for parallel networks. Each router or networking device used on the path between the two end resources requiring RSVP participates in delivering the QoS demanded by the multimedia application.

Network congestion is monitored and managed through the implementation of Random Early Detection (RED). During peak traffic loads, transmission volume can lead to network congestion. RED works in concert with RSVP to maintain end-to-end QoS during these peak loads by selectively dropping traffic at the source using TCP slowstart characteristics. Thus, the source stations feeding into the network slow down their feed until the network metrics defined for the low watermark against RED are met.

Generic traffic shaping works in a similar fashion to RED. However, generic traffic shaping, also called interface-independent traffic shaping, reduces the flow of outbound traffic to the network backbone. This takes effect when a router connecting to a network backbone composed of Frame Relay, SMDS, or Ethernet receives Layer 2-type congestion packets from downstream network transport devices. Generic traffic shaping throttles back the outbound traffic entering the backbone network at the source of entry.

Secure Data Transmission

Security, privacy, and confidentiality over public or untrusted IP networks are paramount for using Virtual Private Networks (VPN). Cisco IOS Release 11.2 reduces the exposure by enabling the capability to provide router authentication and network-layer encryption. Router authentication enables two routers to exchange a two-way Digital Signature Standard (DSS) public key before transmitting encrypted traffic over VPNs using generic routing encapsulation (GRE). The exchange is performed once to authenticate the routers by comparing the hash signature of the keys.

Network-layer encryption uses Diffie-Hellman keys for security. These keys form a Data Encryption Standard (DES) 40- or 56-bit session key. The keys are configurable and set a "crypto-map" that use extended IP access lists to define network, subnet, host, and/or protocol pairs requiring encryption between routers.

Support for IBM Networking Environments

Cisco has been the leader in providing SNA and NetBIOS support over IP networks. Cisco IOS has several means for transporting IBM type traffic, specifically SNA, over router backbone networks. The basis for the transport is encapsulation. Cisco IOS has five different encapsulation techniques and supports full APPN functionality in its native form. The five encapsulation techniques are the following:

- Remote Source Route Bridging (RSRB)
- Serial Tunneling (STUN)
- Data Link Switching Plus (DLSw+)
- Frame Relay RFC 1490
- Native Client Interface Architecture (NCIA)

Along with the five encapsulation techniques, Cisco IOS supports SDLC-to-LLC2 (SDLLC) conversion. This allows SNA devices using the IBM SDLC protocol to attach serially to the router, as if the router were functioning as an IBM front-end processor. SDLLC converts the SDLC frame into a LLC2 frame for transmission to SNA PU devices connected to Cisco routers by either serial, frame relay, or LAN.

IBM configuration and connectivity are also enhanced using Cisco IOS as a TN3270 server. TN3270 is an IETF RC standard that allows non-SNA devices to act as IBM 3270 terminals. Routers using Cisco IOS can act as a TN3270 server for these devices and present their representation to the mainframe as IBM 3270 terminals attached to IBM 3174 control units.

Direct connectivity to the mainframe from a Cisco router is accomplished using a Channel Interface Processor (CIP). The CIP can connect the Cisco 7x00 router series to the mainframe using ESCON or block multiplexing channel connectivity. The CIP provides for SNA TCP/IP services for connecting to the mainframe. The 7200 series uses a Channel Processor Adapter (CPA), which has all the same functions and features as the CIP.

Two management enhancements for supporting IBM SNA over Cisco routers enable SNA network management and performance. Cisco IOS now supports the IBM NMVT command, which is set for sending alerts to the mainframe network management system (NetView) when SNA devices defined to the router have outages or errors. The IOS also has a Response Time Reporter (RTR) feature, allowing operators to analyze SNA response time problems on each leg of the path to the mainframe form of the end user device. This is extremely important to determine bottlenecks in the Cisco router network affecting SNA response time problems.

IP Routing Protocols

Cisco IOS supports a variety of routing protocols. Two of these are Cisco developed and therefore are considered proprietary. All other routing protocols are international standards. The two Cisco routing protocols are Interior Gateway Protocol (IGRP) and Enhanced (IGRP).

IGRP supports IP and ISO CLNS networks. IGRP has its roots in distance-vector-transport-routing schemas with enhancements for determining the best route based on a four-part composite metric that includes bandwidth, delay, load, and reliability. In this decision process, IGRP assumes that the route with the highest aggregate should be the preferred route. However, it does not take into account bandwidth utilization and can therefore itself overload a route and cause congestion. Enhanced IGRP utilizes the Diffusing Update Algorithm (DUAL) along with its roots in link-state-routing protocols to determine the best path between two points. Enhanced IGRP merges the best of distance vector and link-state-routing algorithms to provide greater route decision making control. Enhanced IGRP has support for routing IP, AppleTalk, and IPX natively.

The following list provides the remaining open standard routing protocols available for use on Cisco routers:

- Routing Information Protocol (RIP)
- RIP2
- Exterior Gateway Protocol (EGP)
- Border Gateway Protocol (BGP)
- BGP4
- Protocol Independent Multicast (PIM)
- Intermediate System-Intermediate System (IS-IS)
- Next Hop Routing Protocol (NHRP)

Bridging

Independent Local Area Networks (LANs) have traditionally been bridged together to expand their size and reach. There are two bridging techniques that all others are based on: Transparent and Source Route. Transparent bridging is also known as a learning bridge. This type of bridge is the type typically found bridging Ethernet LANs. Cisco IOS supports the following Transparent bridging features:

- IEEE 802.1(d) Spanning-Tree Protocol
- IEEE 802.10 virtual LANs
- DEC spanning tree
- Bridging over X.25 and Frame Relay networks
- Remote bridging over synchronous serial lines

Source Route bridging provides the path between session partners within the frame itself. Transparent bridging has been coupled with Source Route bridging to allow both techniques to be operable on the same interface. This bridging technique is known as Source Route Transparent (SRT) bridging. Another type of bridging that enables the passing of LAN frames from an Ethernet to a Token Ring LAN is called Source Route/Translational Bridging (SR/TLB). This bridging technique enables SNA devices on an Ethernet, for example, to communicate with the mainframe off of a Token ring LAN.

Packet Switching

Packet switching has its foundation in X.25 networks. Today the most widespread use of packet switching is considered to be frame relay. Cisco provides packet switching for frame relay, SMDS, and X.25 for corporate network support. The most comprehensive of these is frame relay. Cisco IOS supports the following functions and enhancements to frame relay networking:

- Virtual subinterface
- TCP/IP header compression
- Broadcast queue
- Frame Relay switching
- RFC 1490—multiprotocol encapsulation
- RFC 1293—Frame Relay Inverse ARP for IP, IPX, AppleTalk, and DECnet
- Discard eligible (DE) or tagged traffic bit support
- LMI, ANSI Annex D, and CCITT Annex A support
- Dial backup
- Frame Relay over ISDN
- Autoinstall over Frame Relay
- RFC 1490—Transparent bridging
- Frame Relay dial backup per DLCI
- Fast Switched Frame Relay bridging
- DLCI Prioritization
- Frame Relay Switched Virtual Circuit (SVC) support
- Dynamic modification of network topologies with any-to-any connectivity
- Dynamic network bandwidth allocation or bandwidth-on-demand
- Backup for PVC backbones
- Resources allocated only when the connection is required to transfer data in private networks
- Traffic shaping over Frame Relay
- Rate enforcement on a per-VC basis
- Per-VC backward explicit congestion notification (BECN) support
- VC-level priority/custom/weighted-fair queuing (PQ/CQ/WFQ) support

NetFlow Switching

Details of session flows through the router network used to be an elusive quest for the network management team. Cisco IOS NetFlow Switching provides “call detail recording” of traffic through the network on both the network and transport layers. This allows Cisco IOS to manage traffic on a per-user, per-application basis. It does this using a connection-oriented model of the end-to-end flows, applying relevant services to the flow of data. What makes NetFlow even more attainable, it is accomplished in software without added hardware features on the Cisco 7500 and 7000 series routers using Route Switch Processor (RSP) or Versatile Interface Processor (VIP) boards.

ATM

Cisco IOS is fully compliant with all the ATM standards. Cisco itself is very active in establishing the ATM standards and, as such, has a complete feature set. Cisco IOS supports all the ATM standards, including the following:

- ATM point-to-multipoint signaling
- ATM Interim Local Management Interface (ILMI)
- RFC 1577—classical IP and ARP over ATM
- SVC Idle Disconnect
- Bridged ELANs
- LANE (LAN Emulation) MIBs
- SSRP (Simple Server Redundancy Protocol) for LANE
- HSRP for LANE
- DECnet-routing support for LANE
- UNI 3.1 signaling
- Rate queues for SVCs per subinterface
- AToM MIB

Dial-on-demand Routing

As mentioned earlier, Cisco supports dial-on-demand (DDR) services that enhance the availability and performance of internetworks. DDR uses switched circuit connections through public telephone networks. Using

these switched circuits allows Cisco routers to provide reliable backup and bandwidth optimization between locations. The features supported by Cisco DDR include

- POTS via an external modem
- SW56 via an external CSU
- ISDN (BRI and PRI) via integrated ISDN interfaces or external terminal adapters
- Dial backup
- Supplementary bandwidth
- Bandwidth-on-demand
- Snapshot routing
- Multiprotocol routing and transparent bridging over switched circuits
- ISDN fast switching
- Asynchronous ISDN access

Access Server

Cisco routers that function primarily as devices for remote users to access the network are referred to as access servers. These access servers support all the features of DDR with enhancements to support terminal types, connection protocols, security, management, and virtual private networks over the Internet. Access servers provide the following services and features:

- Asynchronous terminal services that include X.25 packet assembler/disassembler (PAD), TN3270, Telnet, and rlogin
- Remote node access over a telephone network using Point-to-Point Protocol (PPP, IPCP, and IPXCP), Xremote, SLIP, and compressed SLIP (CSLIP), AppleTalk Remote Access (ARA) protocol versions 1 and 2, and MacIP
- Multichassis Multilink PPP (MMP)—an aggregate methodology for sharing B channels transparently across multiple routers or access servers
- Asynchronous routing—IP, IPX, and AppleTalk routing
- TN3270 enhancements
- PPP/SLIP on protocol translator virtual terminals
- TACACS+

- TACACS+ single connection
- TACACS+ SENDAUTH function
- ATCP for PPP
- Asynchronous mobility, which connects users to private networks through public networks such as the Internet
- Asynchronous callback, in which a router recognizes a callback request and initiates the callback to the caller
- Asynchronous master interfaces, which are templates of standard interface configurations for multiple asynchronous interfaces on the access server
- ARAP and IPX on virtual asynchronous interfaces
- Local IP Pooling, a pool of reusable IP addresses assigned arbitrarily to asynchronous interfaces
- Remote node NetBEUI, which uses PPP Network Control Protocol (NCP) for NetBEUI over PPP NetBIOS Frames Control Protocol (NBFCP)
- Modem auto-configuring, which is auto-discovery and auto-identification of attached modems, allowing for automatic modem configuration
- NASI (Novell Asynchronous Services Interface)
- RFC 1413 Ident
- RADIUS (Remote Authentication Dial-In User Service)
- Virtual Private Dial-up Network (VPDN)
- Dialer profiles
- Combinet Packet Protocol (CPP)
- Half bridge/half router for CPP and PPP

LAN Extension

Cisco central-site routers, like the 7x00 series, can extend their LAN connectivity over a WAN link using Cisco IOS LAN Extension. The central site router configures LAN Extension services to a multilayer switch at the remote site in a hub-and-spoke configuration. This connection provides a logical extension of the central sites LAN to the remote.

The LAN extension is a practical use of Cisco's CiscoFusion architecture. CiscoFusion describes the combined use of Layer 2 switching or bridging with Layer 3 switching or routing. This combination provides transparent connectivity under the LAN extension supporting



- IP
- IPX
- AppleTalk
- DECnet
- VINES
- XNS protocols.

Since LAN extension supports functions of Layer 2 and 3, MAC address and protocol filtering and priority queuing are accomplished over the WAN links for efficient use of bandwidth.

CHAPTER

2

Cisco Router Hardware



The Cisco router product line has three flavors. Cisco routers are available as modular, fixed, or combination configurations. Along with full-router configuration, Cisco offers router platforms on personal computer (PC) card formats. Additionally, Cisco combines routers and small hubs into one device suitable for small office installations. The key to a successful implementation of Cisco routers in a networking environment is proper placement and configuration of the router. Each Cisco router offering is suited for a specific function. These functions are depicted in Figure 2-1 as core, distribution, and access. These functional characteristics make up Cisco's router inter-network architecture.

Cisco Router Network Hierarchy

Early on in the development of internetworks, an architecture emerged. This architecture for deploying routers was incorporated into a hierarchical architecture that Cisco employs and preaches to its customer base. The architecture relies on the capability of the processor in the router and its need for processing routes, filters, and physical connections. It places the larger Cisco 7x00 series and 12000 series routers at the center or core of the network. The 4x00 series routers are at the net layer of the network architecture called the distribution layer. Finally, the 25xx, 100x, and 7x0 series routers constitute the access layer of the architecture.

Although these assignments to the three different layers of the architecture make sense, it does not mean that 7x00 series routers cannot be used as a distribution or access router. Likewise, in some cases, the 4500 and 4700 series router platforms can be used as a core or access router. However, the smaller fixed and combination routers are most suited for the access layer and will not perform the physical or logical requirements of the core or distribution routers.

Core

The routers that comprise the core layer of the architecture are often referred to as the backbone routers. These routers connect to other core routers, providing multiple paths over the backbone between destinations. These routers carry the bulk of WAN traffic between the distribution routers. Core routers are usually configured with several high-speed interfaces, as shown in Figure 2-2.

Figure 2-1
Cisco router
hierarchical network.





Figure 2-2
Core router
bandwidth and line
network
configuration.



However, due to the introduction of ATM and interface cards that provide up to OC-12 speeds (622 Mbps), core routers may only require two physical interfaces. However, as the section on ATM configuration will reveal, multiple subinterfaces are allowed on each physical interface. The need for the core router to manage many high-speed interfaces is still a requirement, even with only two physical ATM interfaces.

The use of Packet over SONET is another alternative to providing a high-speed core using Cisco routers. In large wide area networks (WANs) and metropolitan area networks (MANs), it is common to have the backbone built on synchronous optical network (SONET) rings with OC-3, OC-12, and OC-48 connections. Packet over SONET allows for the transmission of IP directly over the SONET network without the use of ATM. This provides a great incentive to corporations that have yet to embrace ATM but have a need for high speed and bandwidth over their backbone. Using Packet over SONET as the backbone transport requires an investment in only routers versus ATM, which requires investments in routers and switches.

Distribution

The distribution router functions as the main conduit for a location back to the core. As an example, in Figure 2-3, the distribution router acts as a core router for a campus environment but as a distribution router for a building. Or the distribution router may act solely as a distribution router for a region or campus, managing only the transmission of data between the core and the access layers.

Access

The outer layer of the architecture is the access layer. It is at this layer that end users gain access to the network resources connected by the routers. A typical example for using access routers is in large buildings or campuses. As depicted in Figure 2-4, access routers connect workgroups and/or floor segments within a building to the distribution router. Access routers also provide remote dial-up connectivity for temporary connections.

Figure 2-3
Distribution router network configuration.





Figure 2-4
Access router network configuration.



Online Insertion and Removal (OIR)

Many networks require 24x7 uptime. Powering down a router to replace or add new interface cards causes an outage to all the LAN segments and WAN connections. Cisco IOS, along with the hardware, has implemented a technique to avoid unnecessary downtime called Online Insertion and Removal (OIR).

Supported Platforms

OIR is specific to the high-end router platforms. The Cisco 7000, 7200, 7500, and 12000 series routers all support the OIR feature. The OIR feature works with all interface processor boards, allowing the router power and non-affected interface cards to remain online and functional.

OIR Process

Removing an interface processor board can be accomplished at any time. A new interface processor board is then installed in the available slot and the route processor will recognize that a new board has been installed. If the newly installed board is a higher density or replacement board with equivalent interfaces (such as Ethernet), the processor board recognizes that the boards are similar in function and automatically configures the interfaces to reflect the previous board's configuration. In this way, OIR reduces operator intervention, thereby eliminating configuration input errors on the new interface processor board.

Exceptions to Using OIR

OIR is specific to interface processors for all interface types and does not support the dynamic replacement of a route processor, route switch processor, or a network engine processor. Replacing these boards requires that the router be powered off. However, if you are using the 7507 or 7513 series routers and have taken advantage of the High System Availability (HSA) feature with Route Switch Processors 2 or 4 (RSP2 or RSP4) OIR removes this restriction. HSA enables these router platforms to operate with two

RSP boards. By default, the RSP installed in the first RSP slot is the system master and the second RSP slot is the system slave. Using HSA, it is now possible to remove an RSP for upgrading or for replacement without disrupting the power to the router or interrupting processing the interface processors.

Cisco 12000 Series

The 12000 series router platform is built in support of providing gigabit (Gb) speeds across WAN and MAN backbones. The Cisco 12000 series is targeted at scaling Internet and enterprise backbones at speeds up to two to four Gbps. This is the aggregate bandwidth of an OC-48 SONET connection. The Cisco 12000 series is optimized for IP only networks and thereby provides a high-speed backbone infrastructure for IP-based networks. The capability to handle OC-3 through OC-48 SONET connections enables network engineers to expand the backbone switching capacity with a range from five to 60 Gbps. Since the 12000 series is built for providing a core backbone, it is designed for maximum uptime and minimal disruption. These features are found in its architecture for

- Redundant switch fabric design
- Line card redundancy
- Dual gigabit route processors
- Online software configuration

The speeds of the Cisco 12000 series routers are possible from the synchronized circuitry of two cards: the Clock and Scheduler card (CSC) and the Switch Fabric card (SFC). Both the CSC and SFC provide an OC-12 switching bandwidth between the line cards for the system. Each type of card has a switching capacity of 15 Gbps.

A minimum of one CSC is required in the router. The CSC performs the following functions for the router:

- System Clock: Clicking is sent to all line cards, GRPs, and SFCs. It synchronizes data transfers between the various components of the system. In redundant mode, the CSC clocks are synchronized for failover.
- Schedule: The scheduler function handles requests from the line cards and schedules when the line card can have access to the switch fabric.

The SFC provides the following functionality for the router:

- Contains only switching fabric.
- Carries traffic between line cards and GRP.
- Receives scheduling and clocking from the CSC.

The chassis configuration of the Cisco 12000 router comes with an upper and lower cage. The upper cage is used mostly for the line cards to connect to the network in addition to the Gigabit Route Processor (GRP) card. The lower cage supplements the ability for the 12000 series router to perform switching by having extra slots for the SFC installs. For more information on the specific cage configurations of the 12000 series router, consult the section specific to the model.

The 12000 series comes in three models: 12004, 12008 and 12012.

Cisco 12004 Series



The Cisco 12004 series is the smallest of the 12000 line. It provides a total of four interface slots and two slots for GRPs. The 12004 supports all the available interfaces of the 12000 series. It is usually used in IP SONET backbone networks with minimal connectivity requirements and is typically used for OC-3 and OC-12 interface connections. The 12004 has an IP datagram switching capacity of five Gbps.

In a single CSC configuration, the 12004 supports OC-12 data rates and a 1.25-Gbps switching capacity. Using redundant CSCs in the two center slots of the upper cage and three SFCs in the lower cage, the 12004 can support OC-48 data rates with a switching capacity of five Gbps. In a redundant GRP configuration, the 12004 has two line card slots available for network connectivity.

Cisco 12008 Series

The Cisco 12008 can switch IP data grams in the range of 10 to 40 Gbps. The minimum configuration requirement for the Cisco 12008 is the presence of a single GRP and a single CSC. As shown in Figure 2-5, the CSC must be placed in either of the two center slots in the upper cage of the 12008. A second CSC can then be placed in the open CSC slot for redundancy. The GRP can be placed in any of the remaining slots. A second GRP can be installed for redundancy in any of the remaining slots. Using redundant GRPs leaves six available slots for line card connectivity to the network. The lower cage houses the three optional slots for use by SFCs.

Figure 2-5
Cisco 12008 slot
configuration.



Installation of a second CSC does not increase the switching capacity but provides redundancy. The addition of the three SFCs enables the router to move from an OC-12 with a switching capacity of 10 Gbps to support an OC-48 data rate with a switching capacity to 40 Gbps with full redundancy, should either the CSC fail or a single SFC fail.

Cisco 12012 Series

The Cisco 12012 has the capacity to switch IP datagrams anywhere from 15 to 60 Gbps. The increase in interface density of the 12012 is created by expanding the lower cage. The lower cage of the 12012 contains five keyed slots, shown in Figure 2-6, for placing the CSC in slots 0 or 1 and the SFCs in slots 2 through 4. The GRP is still installed in the upper cage. In a redundant GRP configuration, there are 10 open line card slots for network connections. The single CSC configuration supports an OC-12 data rate and a capacity of 15 Gbps switching. A redundant CSC configuration with three SFCs installed enables the 12012 to support OC-48 data rates and a switching capacity of 60 Gbps.

12000 Usage


The 12000 series is placed at the very core of the network. Since it is optimized for IP traffic, the network must be designed so that IP traffic only flows through these routers. For example, in a network that is based on IP and SNA, the SNA data must be transported using RSRB or DLSw+ with TCP or FST encapsulation techniques. In this manner, the high-speed backbone can be used for connecting remote locations to the main data centers. Likewise, using Voice over IP, the router or PBX must encapsulate the voice data into IP prior to delivering it to the 12000 series backbone routers. Based on this type of usage, the 12000 series is ideal for

- Internet service providers (ISPs)
- Carriers providing Internet services and utilities
- Competitive access providers (CAPs)
- Enterprise WAN backbones
- MAN backbones


Figure 2-6

Cisco 12012 slot configuration.

Rearview of Cisco 12012 GSR




Up to 11 slots are available for line cards



12000 Switch Processors

The Cisco 12000 Gigabit Route Processor is based on the IDT R5000 Reduced Instruction Set Computer (RISC) CPU. This processor, shown in Figure 2-7, has an external bus clock speed of 100 MHz and an internal clock speed of 200 MHz. All the models of the Cisco 12000 series routers use the same GRP card. The GRP can be installed in any slot of the 12012 except for the far right slot, which is reserved for the alarm card. Normal practice is to install the first GRP in the far left slot. On the 12008, the GRP can be installed in any available slot of the upper cage except for the two center slots, which are reserved for the Clock and Scheduler Cards.

Figure 2-7
Gigabit Router
Processor (GRP).



12000 Memory

Each GRP comes with a base of 64 MB of dynamic random-access memory (DRAM), which is upgradeable to 256 MB of parity-protected extended data output (EDO) DRAM. The DRAM is provided in two dual in-line memory module (DIMM) formats running at 60 nanoseconds (ns). The GRP uses the DRAM for storing systems software (Cisco IOS), configuration files, and line card routing tables. The Cisco IOS runs from DRAM. Table 2-1 lists the DRAM socket locations and DRAM configurations for upgrading from 64 to 256 MB.

TABLE 2-1

12000 Series
DRAM Update
Configurations

| Total DRAM | DRAM Socket | Number of DIMMs |
|------------|-------------------------------|-----------------|
| 64 MB | U39 (Bank 1) | 1 (64 MB DIMM) |
| 128 MB | U39 (Bank 1) and U42 (Bank 2) | 2 (64 MB DIMM) |
| 128 MB | U39 (Bank 1) | 1 (128 MB DIMM) |
| 256 MB | U39 (Bank 1) and U42 (Bank 2) | 2 (128 MB DIMM) |

In addition to DRAM, the GRP also includes Static RAM (SRAM) and Non-volatile RAM (NVRAM). The SRAM provides 512 KB of secondary CPU cache memory functions. The SRAM cannot be configured by the user, nor can it be upgraded in the field. The SRAM is primarily a staging area for routing table updates to and from the line cards. The NVRAM stores router configurations, system cache information, and read-only memory (ROM) monitor variables in 512 KB. Information stored in NVRAM is available even after the router loses power, while SRAM and DRAM lose the information stored within them. Like SRAM, the NVRAM cannot be configured by the user, nor can it be upgraded.

The GRP also utilizes flash memory. There is 8 MB of single inline memory modules (SIMM) on the GRP for storing Cisco IOS software images as well as saving router configurations and other types of end user files. Additionally, the only board flash memory can be coupled with the capability of using 20 MB PCMCIA flash memory cards that install on two slots on the GRP with a total capacity of 40 MB. Each card can be used for storing Cisco IOS software images and other files required by the router for operation.

For operational support, the GRP enables remote access to the Cisco 12000 router through an auxiliary dial-up port in an IEEE 802.3 10/100

Mbps Ethernet port for Telnet connections. In addition, the GRP has an RS-232 console port connection for direct serial connectivity from a PC to the router.

The GRP can be installed in any of the slots available in the upper cage of the Cisco 12000 series routers. The exception to this is the Cisco 12012 where the GRP cannot be installed in the far right slot. This slot is reserved for the alarm card.

12000 Line Cards

Each line card is comprised of several functions equivalent on each card. The line card uses burst buffers to prevent packet dropping when there is an instantaneous increase in back-to-back small packets queued for transmission. Burst buffers increase throughput and maintain an even packet burst for packets arriving on Layer 3 switch processing.

Each line card contains two silicon queuing engines, one for receive and one for transmit. The receiving engine moves packets from burst buffers to the switch fabric. The transmit moves the packets from the switch fabric to the transmit interface. The silicon engine also manages the movement of IP packets in buffer memory, which defaults to 32 MB, split evenly between receive and transmit buffers. The amount of buffer memory in use is configurable up to 64 MB for receive and 64 MB for transmit.

An application-specific integrated circuit (ASIC) is used for supporting the high-speed process required to perform Layer 2 switching. To assist in the decision making, an IDT R5000 200 MHz RISC processor is on the line card to make forwarding decisions based on the Cisco Express Forwarding table and the Layer 2 and Layer 3 information in the packet. The GRP is constantly updating the table based on information gathered from the routing table.

The line card also contains a switch fabric interface. This is the same 1.25 Gbps full-duplex data path used by the GRP. When a packet is on the proper queue, the switch fabric requests the CSC for scheduling the transfer of the packet across the switching fabric.

Also, a maintenance bus module on the line card provides the master Mbus module of the GRP with requested information, which is reported in temperature and voltage. In addition, the Mbus on the line card stores the serial number, hardware revision level, and other pertinent information about the card in EEPROM.

Each line card maintains the Cisco Express Forwarding (CEF) table. The table is built on routing table information provided by the GRP and is used to make forwarding decisions.

Six line cards are available for connecting the 12000 series router to the network:

- Quad OC-3c/STM-1c Packet-Over-SONET (POS)
- Quad OC-3 ATM line card
- OC-12c/STM-4c Packet-Over-SONET (POS)
- OC-12c/STM-4c Asynchronous Transfer Mode (ATM)
- OC-48c/STM 16 Optical IP Interface card
- Channelized OC-12 Line card

The Quad OC-3c/STM-1c Packet-Over-SONET (POS) is shown in Figure 2-8.

The card has four ports for interfacing directly to the SONET provider's equipment. The Quad OC-3c/STM-1c Packet-Over-SONET (POS) line card must be ordered for either single mode or multimode SC fiber connections. Each mode supports full-duplex transmission. The card uses 128 KB burst buffers to prevent packet dropping when there is an instantaneous increase in back-to-back small packets queued for transmission.

The Quad OC-3 ATM Line card shown in Figure 2-9 performs ATM segmentation and reassembly functions for ATM connectivity. Segmentation is the process of converting packets to ATM cells, while reassembly is the process of converting ATM cells to packets. The Quad OC-3 ATM Line card can handle up to 4,000 simultaneous reassemblies of an average packet size of 280 bytes. To perform this capability, the segmentation and reassembly is performed on ASIC. The ASICs also allow each of the four ports on the Quad OC-3 ATM Line card to support 2,000 active virtual circuits. The card must be ordered as either single mode or multimode fiber connection. The Quad OC-3 ATM Line card supports a burst buffer of 4 MB.

The OC-12c/STM-4c Packet-Over-SONET (POS) illustrated in Figure 2-10 has a one duplex SC single- or multimode fiber connection. The port supports OC-12c at a 622 Mbps data rate. The OC-12c/STM-4c Packet-Over-SONET (POS) has a burst buffer of 512 KB.

The OC-48c/STM 16 Optical IP Interface card is shown in Figure 2-11, a single duplex SC or FC single-mode fiber connection. The top port is the transmit (TX) connection and the bottom port is the receive (RX) connection. The interface supports a full 2.5 Gbps optimized for transporting packet over SONET (POS). The burst buffer on the OC-48c/STM-16 Optical Interface card is 512 KB with a default buffer memory of 32 MB for receiving and 32 MB for transmitting. Cisco IOS software Release 11.2(14) GS1 and line card microcode Version 1.14 is required for complete support of all features. The typical maximum distance the line card can sustain is 1.2 miles or two kilometers.

Figure 2-8
The Quad OC-3c/ATM-1c Packet-Over-SONET (POS) Interface card.




Figure 2-9
The Quad OC-3 ATM
Line card.




Figure 2-10
The OC-12c/STM-4c
Packet-Over-SONET
(POS) Interface card.





Figure 2-11
The OC-48c/STM 16
Optical IP Interface
card.




The Channelized OC-12 Line card shown in Figure 2-12 supports only single mode full-duplex SC connections at 622 Mbps. Its burst buffer size is 512 KB. The forwarding processor on the Channelized OC-12 Line card is an IDT R5000 RISC processor rated at 250 MHz.

12000 Software Support

The Cisco IOS software for the Cisco 12000 series routers is optimized for transporting IP traffic. The first release of Cisco IOS supporting the Cisco 12000 series platform is the 11.2 release. The Cisco IOS Release 11.2 and higher supports the following IP IOS functions:

- Routing protocols
 - Interior: RIP, OSPF, IS-IS, ISO/CLNP, EIGRP, EGP
 - Exterior: BGP
- Routed protocols
 - IP
- BGP4 support
 - Route reflections
 - MED (Multi-Exit Discriminators)
 - Communities
 - DPA (Destination Preference Attribute)
 - Flat/weighted route dampening
 - Confederations
 - Next hop-self
 - EGP multipath
 - Static routing (IGP)
- Management and access
 - SNMP,
 - Telnet,
 - MIB II

Figure 2-12
The Channelized OC-12 Line card.




11704

Cisco 7500 Series

The Cisco 7500 series router is the high-end routing platform for supporting corporate enterprise-wide networks as well as a keystone for the Internet backbone itself. The port capacity and available interface types enable the 7500 to serve all layers of Cisco's routing architecture. The speed with which the 7500 series processes packets between the various interfaces is due to the use of high-speed bus architectures. The architecture is called the Cisco Extended Bus (CyBus). The CyBus supports any combination of interface processors on the 7500 series platform. The CyBus has an aggregate throughput of 1.067 Gbps. The 7500 series encompasses three models: Cisco

Figure 2-13

The Cisco 7505 series router.



7505, Cisco 7507, and the high-end of the platform, Cisco 7513. Each model has a specific location for the RSP boards. The 7500 series platform supports 15 different feature sets.

Cisco 7505 Series

The 7505 series is the smallest platform of the 7500 line. It supports four interface processors and one RSP board. Figure 2-13 depicts the platform format for the 7505. The 7505 comes with a single CyBus for attaching the interface boards to the RSP. The 7505 series supports RSP1 and RSP4. The single power supply offered on this platform makes the 7505 series a choice for locations with low availability requirements, but with high throughput requirements and the need for varied interface support.

Cisco 7507 Series

The Cisco 7507 series router platform from Cisco expands the interface combination possibilities by providing five slots for interface processors, as shown in Figure 2-14. The 7507 series provides a higher reliability through the use of a second power supply and dual RSP boards. The redundant configuration for the 7507 series enables it to reliably serve as a core or distribution router. The 7507 series uses either an RSP2 or RSP4. The RSPs used in a dual RSP configuration (HSA) should, however, be the same RSP platform.

Added to the higher availability architecture of the 7507 is the use of a dual CyBus architecture. This architecture not only enables recovery

should a bus fail, but the architecture allows both buses to be used simultaneously, allowing higher throughput than on the 7505 series.

Cisco 7513 Series

The Cisco 7513 is the high-capacity 7500 series router platform from Cisco. This series provides two RSP slots for HSA and 11 interface processor slots, as shown in Figure 2-15. These support any combination of network interface requirements.

The 7513 series also supports the dual CyBus architecture and allows for two power supplies. Both RSP2 and RSP4 processors are supported on the platform. The 7513's high capacity for interfaces makes it a useful platform for multiple LAN segment interfaces in a large environment, along with using the interface combination possibilities to serve as a core, distribution, or access router.

Figure 2-14
The Cisco 7507 series router.





Figure 2-15
The Cisco 7513 series
router.



7500 Usage

The 7500 series is quite versatile and provides the functionality of core, distribution, and access layers. The 7505 is used as a low-availability access router servicing a casual end user site supporting multiple LAN interfaces. A site of this nature is usually autonomous with processing done locally for the majority of the time.

The 7507 series servicing the remotes performs the functions of the distribution and access layers. The 7507 features are useful in access locations where there are many different types of interface requirements and many LAN segments. The features also support a high volume of data from the site to the WAN. As a WAN distribution router, the 7507 connects many of the remote access locations without going to the core routers. The 7513, as indicated earlier, is suitable for all three layers of the router networking architecture.

7500 System Processors

The Route Switch Processor (RSP) platform used on the 7500 series router is a combination of the router processor (RP) and switch processor (SP) originally used on the Cisco 7000 series router platform. Combining the functionality of the RP and SP into one board enables the RSP to switch and process packets faster and allows each platform to gain an extra slot for an interface processor.

Three types of RSP platforms exist in the 7500 series. The base platform of each RSP type comes with 32 MB of DRAM and 8 MB of Flash SIMM memory. The 7500 series uses the Flash SIMM for storing and loading the Cisco IOS BOOT images necessary for the RSP to activate prior to executing any other functions. The DRAM is upgradeable from 32, 64, or 128 MB of DRAM with Flash memory upgrades using PCMCIA cards in up to two slots totaling 40 MB. Each RSP comes with 128 KB of NVRAM to store the IOS system running and startup configuration files.

RSP1 The RSP1 is the default RSP on the 7505 series router and is available only on this router. The RSP1 stores the Cisco IOS image in Flash memory on the RSP or on up to two Intel Series 2+ Flash memory PCMCIA cards. The RSP1 has an external clock speed (bus speed) of 50 MHz and an internal clock speed (CPU speed) of 100 MHz.

RSP2 The RSP2 is the base RSP board supplied for the 7507 and 7513 series routers. The RSP2 operates at an external clock speed (bus speed) of up to 50 MHz and an internal clock speed (CPU speed) of 100 MHz. The RSP2 platform of the RSP system processors supports the High System Availability (HSA) features. Using two RSP2 system processors, the 7507 and 7513 provide for RSP failure recovery, as the slave takes over for the

master if the master should experience an outage. The default for identifying the system master is the RSP2 occupying slot 2 on the 7507 and slot 6 on the 7513 router. The order is configurable, but it is highly recommended that the defaults be taken when using HSA. A caveat to using HSA is Cisco IOS Release 11.1(5) or higher and ROM monitor version 11.1(2) or higher. Each RSP2 must have the same version of ROM monitor installed for HSA to function properly.

RSP4 The RSP4 platform of the RSP system processors is available for the three 7500 series platforms. Its external clocking speed (bus speed) is 100 MHz and supports an internal clocking speed (CPU speed) of 200 MHz. The RSP4 also uses DIMM chip sets for DRAM memory. As such, the RSP4 DRAM configuration is 32, 64, 128, or 256 MB.

An enhancement to the RSP4 over the RSP1 and RSP2 is the use of static RAM (SRAM) for packet buffering and a secondary cache memory for CPU functions. The RSP4 supports any type of PCMCIA flash memory card for flash memory.

PCMCIA card formats come in three types: PCMCIA Type 1 and 2 are usable in slot 0 and slot 1. Type 3 PCMCIA flash memory cards are only supported in slot 1 of the PCMCIA slots for the RSP4. Like the RSP2, the RSP4 supports HSA. Support for HSA on the RSP4 is dependent to the level of Cisco IOS and ROM monitor. HSA is fully supported on the RSP4, using Cisco IOS release 11.1(8)CA1 and ROM monitor version 11.1(8)CA1 and higher.

7500 Memory

Memory on the RSP and any interface processor is paramount to efficiently running the routers. The more, the better. It does not hurt to order the highest amount of memory available for any platform as an inexpensive insurance policy against poor design or "memory leaks" from the IOS or microcode software.

That aside, the 7500 series platform comes with DRAM memory size recommendations based on the number of IP routes in a network. Cisco categorizes network sizes into the following:

- Small networks: less than 2,000 IP routes
- Medium networks: between 2,000 and 10,000 IP routes
- Large networks: greater than 10,000 IP routes

The following DRAM memory requirements are recommended for the RSP1, RSP2, and RSP4 system processors on each of the 7505, 7507, and 7513 router platforms:

- Small networks: 32 MB
- Medium networks: 32 MB
- Large networks: 64 MB

Cisco highly recommends that even if some networks are much smaller than the 2,000 IP routes a minimum of 32 MB of DRAM is beneficial for router performance.

The Flash memory PCMCIA cards available for insertion into slot 0 and slot 1 of the RSP boards are available in different memory sizes. The default card comes with 8 MB of memory and has a default IOS software image stored. If a spare is ordered or purchased, it must first be formatted before use. PCMCIA cards used on RP boards from a 7000 series router must be reformatted for use on the 7500 series router due to a difference in the formatting of memory on the different system processors.

7200 Series


The Cisco 7200 series router is a change in the routing platform architecture for Cisco. The architecture of the interface slots is based on the technology conceived with the Versatile Interface Processor 2 (VIP2) boards from the 7x00 series. Instead of using slots, the 7200 series uses port adapters. Figure 2-16 illustrates the adapter layout for the 7200 series router.

The 7200 series platform is available in two formats. The 7204 supports up to four port adapters, while the 7206 supports up to six port adapters. Each platform requires a network processing engine (NPE) and an Input/Output (I/O) Controller processor. The I/O Controller has two slots for PCMCIA Flash memory cards and can be optionally configured with a Fast Ethernet interface using an MII connector. Each port adapter supports the OIR function allowing non-interruption of port upgrades or replacements. As found in the 7x00 series, the replacement of similar adapters is automatically configured on insertion.

The 7200 series uses a peripheral component interconnect (PCI) bus architecture in support of the various network interfaces available using the port adapters. This bus architecture is built on two primary PCI buses and a secondary PCI bus providing a high-speed mid-plane rate of 600

Figure 2-16

The Cisco 7200 series router adapter layout.



Mbps. A second power supply is available for added redundancy enhancing high availability.

7200 Usage

The 7200 is positioned as a low-volume core router or medium distribution router. Network Layer 3 switching support directly supported by the 7200 series makes it an excellent candidate as a distribution router for a large office complex or as an access router for many LAN segments within the office complex.

7200 Network Processing Engine

The maintenance and execution of system management functions are supported by the network processing engine (NPE) on the 7200 series platform. The NPE works with the I/O Controller to monitor environmental and shares in system memory management. Two versions of the NPE exist in the 7200. The NPE-100 maintains an internal clock speed of 100 MHz and an external clock speed of 50 Mhz. The higher performance NPE-150 uses an internal clock speed of 150 MHz and an external clock speed of 75 Mhz. In addition, the NPE-150 includes 1 MB of packet SRAM for storing packets used in fast switching. The NPE requires Cisco IOS software version 11.1(5) or later for the 7206 and 11.1(6) or later for the 7204.

7200 Memory

Memory requirements on the 7200 series are dependent on the varied adapter configurations possible with each platform. The NPE comes standard with 32 MB of DRAM. This memory is incremental in 8, 16, or 32 MB SIMMs totaling 128 MB. Both the NPE-100 and NPE-150 have a unified cache memory of 512 KB as a secondary cache for the Orion R4700 RISC processor. Appendix C details the memory configuration requirements for the 7200 series platforms.

The I/O Controller for the 7200 series provides NVRAM for the storage of system configurations and logging environmental monitor results. The two PCMCIA slots found on the I/O Controller support the Intel Series 2+ Flash Memory PCMCIA formats. These PCMCIA cards have 8, 16, or 20 MB of flash memory on board. The total available for the two slots combined is 40 MB.

7000 Series


The Cisco 7000 series was the original “big” router platform introduced. It was the replacement for the Cisco AGS and AGS+ router platforms. The 7000 platform itself has since been replaced by the 7500 platforms.

The Cisco 7000 comes in two platforms, as Figure 2-17 depicts. These are the 7000 and the 7010 series. The 7000 has a total of seven slots. Five of these slots are used for interface processors and two for system processors. The 7010 series is smaller and offers a total of five slots. Three of the slots on the 7010 are used for interface processors and the remaining two slots provide support for system processors.

OIR was originally introduced with this platform along with a backplane called the Cisco extended bus (CxBus). The CxBus architecture provided a data bus throughput of 533 Mbps on the 7000 series. The 7000 series supports up to two power supplies to enhance availability. However, the series itself does not support the high system availability feature found on the 7500 series platforms.

Figure 2-17

The two platforms of the Cisco 7000 series router.



7000 Usage

The 7000 platforms were initially developed primarily as a core router. However, the need for higher port densities and faster processing have moved the 7000 series out of the core and into the role of a small to medium distribution. The 7000 or 7010 is used as a distribution router servicing access locations.

7000 System Processors

On introduction of the 7000 platform, Cisco used a Motorola 68040 CPU clocked at 25 Mhz. Although this was considered fast for the time, it has since been antiquated. The CPU is found on the Router Processor (RP) board. The RP is installed in slot 6 of the 7000 series and slot 4 of the 7010 series.

In concert with the RP, the 7000 platform utilized three models of a Switch Processor (SP). These consist of the

- Switch Processor (SP)
- Silicon Switch Processor (SSP)
- Silicon Switch Processor-2MB (SSP-2 MB)

The SP offloaded the responsibility of managing the CxBus from the CPU on the RP board, thus allowing the RP to efficiently manage system functions. Further enhancements using a Silicon Switch Engine (SSE) on the SP allowed the SP to examine incoming packet data link and network link header information, making an intelligent decision on whether the packet should be bridged or routed before forwarding the packet to the corresponding interface. The speed of the decision process was enabled by using a silicon-switching cache that kept track of packet information through the router.

The SSE is encoded in the SP hardware and within this configuration is the Silicon Switch Processor (SSP). The SSP performs switching decisions independently of the RP, thereby increasing the throughput and efficiency of system resources. The base SSP includes an extra 512 KB of memory for handling switching decisions, while the SSP-2 MB provides an extra 2 MB of memory. In the 7000 series, the SP, SSP, or SSP-2 MB is installed in slot 5 and in the 7010 series the SP, SSP, or SSP-2 MB is installed in slot 3.

Extending the life of the 7000 platform is made possible by the introduction of the Route Switch Processor 7000 (RSP7000) and the 7000 Chassis Interface (7000CI) processors. These two boards together give the 7000 platform the enhancements and capability to use the IOS software made for the 7500 router platform. The IOS software must be at IOS Version 10.3(9), 11.0(6), 11.1(1) or later to support the RSP7000 processor and the 7000CI processor. The RSP7000 increases the performance of the 7000 platform by using a MIPS Reduced Instruction Set Code (RISC) CPU at 100 MHz and a bus speed clocking (external clock) of 50 Mhz. Use of the RSP7000 on the 7000 and 7010 series routers enables these platforms to use the Versatile Interface Processor (VIP) technology supported under the 7500 IOS software platform. The 7000CI monitors Chassis-specific functions, relieving the RSP7000 from the following duties:

- Reporting backplane and arbiter type
- Monitoring power supply status
- Monitoring fan/blower status
- Monitoring temperature sensors on the RSP7000
- Providing router power up/down control
- Providing power supply power-down control

The RSP7000 is installed in slot 5 of the 7000 series and slot 4 of the 7010 series. The 7000CI is installed in slot 6 of the 7000 series and slot 3 of the 7010 series.

7000 Memory

Although both the RP and RSP7000 use the Intel Series 2+ flash memory cards, they must be reformatted if used between the two processors. The RP supports one slot for flash memory and the RSP7000 supports two flash memory slots. The RP flash memory PCMCIA card is either 8 MB or 16 MB. The RSP7000 is available in either 8, 16, or 20 MB formats with a total of 40 MB of flash memory.

The RP processor comes standard with 16 MB of RAM and is upgradeable to 64 MB. The RSP7000 comes standard with 32 MB of RAM with expansion to a total of 128 MB. Appendix D highlights the various DRAM requirements along with the feature sets available for the 7000 series routers.

Cisco 7x00 Series Interface Processors

The strength of the Cisco router product line is the capability to support the many different LAN/WAN physical interface standards available. The Cisco 7x00 family of routers has a very versatile offering supporting these standards without restricting the combinations possible by mixing and matching the interface processor boards on the chassis.

The Cisco 7x00 router platform can actively support any combination of

- Ethernet, Fast Ethernet, Gigabit Ethernet, Token Ring, FDDI, serial, channelized T3, Multichannel E1/T1, IBM mainframe channel attachment, ATM, Packet OC-3, ISDN, and HSSI interfaces.
- These interfaces are provided on interface processors that connect physical networks to the high-speed bus of the Cisco 7x00 router. The interface processors are specific to the 7000 and 7500 router platforms. The 7200 router platform uses port adapters that are akin to the port adapters of the Versatile Interface Processor (VIP) available on the 7000 and 7500 router platforms. The VIP and the port adapters supported are discussed in the following section.


The interface processors are modular circuit boards measuring 11 x 14 inches with network interface connectors. The interface processors all support OIR and are loaded with microcode images bundled with the Cisco IOS software. The exception to this bundling of microcode is the CIP, which is unbundled as of IOS Version 11.1(7) and higher. For the most part, each interface processor is self-contained on a single motherboard. However, some interface processors require a companion board attached to the motherboard. For example, the AIP board uses a physical layer interface module (PLIM), which is installed at the factory based on the AIP order.

7x00 ATM Interface Processor (AIP)

The AIP board supports fiber-optic connectivity and coaxial connectivity in support of Asynchronous Transfer Mode (ATM) networking environments. The board also supports single-mode and multimode fiber-optic connections. Figure 2-18 illustrates the AIP board with a fiber-optic PLIM.

Figure 2-18

The AIP board with a fiber-optic PLIM.



The following lists the media types supported by the AIP board:

- Transparent Asynchronous Transmitter/Receiver Interface (TAXI) multimode fiber-optic
- Synchronous Optical Network (SONET) multimode fiber-optic
- SONET single-mode fiber-optic
- E3 coaxial
- DS3 coaxial

The AIP board can now support up to OC-12 SONET connectivity for high bandwidth and throughput requirements. Each of the media types supported requires a specific cable connection. Appendix E lists all the cable specifications for all the router platforms and their interfaces.

7x00 Channel Interface Processor 2 (CIP2)

The Cisco Channel Interface Processor 2 (CIP2) is the second generation of IBM mainframe channel connectivity boards offered in support of connecting router networks directly to the mainframe. The CIP2 is a direct competitor to IBM's 3172 Interconnect Controller and the IBM 2216 channel-attached router. The CIP2 has memory and processing advantages over the first generation CIP. The CIP2 supports both IBM's parallel bus-and-tag channel and ESCON fiber channel architectures. The CIP2 ships with a default of 32 MB of memory with memory configuration of 64- and 128-MB allocations.

The CIP2 is compatible with the Cisco 7000 series router using Cisco IOS Release 10.2(13) or later, 10.3(12) or later, 11.0(10) or later, and all versions

at 11.1(5) or later. The 7500 series router requires the Cisco IOS release level be at 10.3(13) or later, 11.0(10) or later, and all versions at 11.1(5) or later.


The CIP2 microcode is unbundled from the IOS software as of Release 11.1(7) and must be ordered separately from the IOS when installing a CIP2. The microcode supports the following mainframe connectivity features:

- TCP/IP datagram
- TCP/IP offload
- CIP Systems Network Architecture (CSNA) connectivity using External Communications Adapter (XCA) communications to VTAM
- TN3270 server
- Native Client Interchange Architecture (NCIA) server
- Advanced Peer-to-Peer Network (APPN)

The CIP2 supports different combinations of channel connectivity to the mainframe. These combinations are configured at the factory and must be ordered appropriately. Figure 2-19 diagrams a CIP2 board with a single parallel channel and single ESCON interface configuration. The valid combinations for the CIP2 interfaces are

- Single parallel channel
- Dual parallel channel
- Single ESCON channel
- Dual ESCON channel
- Single ESCON channel and single parallel channel

Figure 2-19
The CIP2 board with a single parallel channel and single ESCON interface configuration.




When ordering a CIP2 board, it is advisable to determine the number of TCP/IP and SNA connections planned for use by the CIP2. The number of connections directly relates to CIP2 performance and memory requirements. Although Cisco has memory recommendations and formulas to calculate memory requirements, it is advisable to order the CIP2 with the maximum amount of memory, 128 MB, to allow for growth and performance without compromising availability and reliability. Appendix E details the CIP2 memory formulas and minimum requirements.

7x00 Channelized T3 Interface Processor (CT3IP)

The CT3IP is based on the VIP2 interface processor architecture. It is a fixed-configuration, meaning that it is not reconfigurable after ordering or installation. The CT3IP supports four T1 connections and a single DS-3 connection, as shown in Figure 2-20.

The T1 connections use a DB-15 connector and the DS-3 uses a transmit (TX) and receive (RX) female BNC connection pair. The DS-3 connection provides up to 28 T1 channels with each channel viewed as a serial interface to the system. Each channel can then be configured individually. The CT3IP board is supported on the Cisco 7500 series and Cisco 7000 series with the RSP7000 and 7000CI boards only.

Figure 2-20
A CT3IP supporting
four T1 connections
and a single DS-3
connection.



7x00 Ethernet Interface Processor (EIP)

The EIP supports 10 Mbps of Ethernet LAN connectivity. Three variations of the EIP board support either two, four, or six 10 Mbps Ethernet 802.3 interface ports. Figure 2-21 diagrams a six-port EIP board.

Attachment of the EIP interfaces may require a transceiver that converts to 802.3 and an attachment user interface (AUI) cable to RJ-45 cable connectivity to a LAN hub or switch.

7x00 Fast Ethernet Interface Processor (FEIP) and FEIP2

The interface processor forms support fast Ethernet connectivity at 100 Mbps. The media supported is twisted-pair or fiber-optic cable. The format of the board uses the port adapter architecture found with VIP2 boards, but the FEIP and FEIP2 port adapters are not interchangeable for use on the VIP2 board or Cisco 7200 series routers. Figure 2-22 illustrates the FEIP and FEIP2 boards. Note that the main difference on the boards is the inclusion of a CPU on the FEIP2. The CPU on the FEIP2 offloads the RSP of switching, filtering, and other previously RSP-based functions, thereby increasing performance on the FEIP2 and the RSP in general.

Both the FEIP and FEIP2 have configurations that support one or two port adapters. Each port adapter supports a RJ-45 and MII connector. The MII connector supports fiber-optic connectivity in concert with a trans-

Figure 2-21
A six-port EIP card.






Figure 2-22
The FEIP and FEIP2 boards.



ceiver. Only one of the interfaces can be active on each port adapter. The RJ-45 supports Category 5 UTP 100BaseTX connectivity, while the FEIP supports full- and half-duplex operations on all interfaces in any combination. The FEIP2 only allows half-duplex operations on the 100BaseTX RJ-45 connection. The FEIP2 can operate both 100BaseFX interfaces using either half-duplex or full-duplex modes. However, in a configuration where both MII interfaces attach 100BaseFX LANs, only one interface can operate in full-duplex mode. In addition to the use of a CPU on the motherboard, the FEIP2 includes 1 MB of SRAM and 8 MB of DRAM.

The Cisco 7000 series supports the FEIP using 100BaseTX with Cisco IOS Release 10.3(5) or later. The Cisco 7500 series supports FEIP 100BaseTX using Cisco IOS software Release 10.3(6) or later. Support for 100BaseFX connectivity on the Cisco 7000 and 7500 series uses Cisco IOS Release 10.3(13) or later, 11.0(10) or later, or Release 11.1(5) or later.

The FEIP2 board and interface support for 100BaseTX and 100BaseFX connections is found in Cisco IOS Release 11.1(10)CA or later for both the Cisco 7000 and 7500 series routers.

7x00 FDDI Interface Processor (FIP)

The FIP enables the Cisco 7000 and 7500 router platform to support single-mode and multimode FDDI connections at 100 Mbps. Figure 2-23 diagrams the four FIP board configurations. These configurations support the following:

- Multimode to multimode with optical bypass
- Multimode to single mode
- Single mode to multimode
- Single mode to single mode with optical bypass

7x00 Fast Serial Interface Processor (FSIP)

The FSIP, as shown in Figure 2-24, uses dual-port port adapters. Each port adapter supports two serial interfaces. Each interface can support up to 6.132 Mbps. The 6.132-Mbps bandwidth is the total allowed for the entire FSIP board. If one or more ports totals a bandwidth of 6.132 Mbps, the remaining ports are not available for use.

Figure 2-23
The FIP board configurations.





Figure 2-24
The FSIP card.




The FSIP supports two configurations: a four-interface serial port adapter and an eight-interface serial port adapter. The first ports are numbered 0 through 3 and the second are numbered 4 through 7.

7x00 High-Speed Serial Interface(HSSI) Interface Processor (HIP)

The HIP is capable of supporting up to 52 Mbps bandwidth. The HIP, diagrammed in Figure 2-25, enables data rates up to 45 Mbps (DS-3) or 34 Mbps (E3) for connecting ATM, SMDS, Frame Relay, or private lines. The HIP uses a special cable and must be ordered from Cisco for supporting this high-speed configuration.

Figure 2-25
The HSSI card.




7x00 Multichannel Interface Processor (MIP)

The MIP, shown in Figure 2-26, is a multichannel multiplexer, allowing the router to emulate an Nx64 or Nx56 backbone multiplexer on a 1.536 Mbps (T1) or 2.048 Mbps (E1) line. The MIP supports seven different types of configurations:

- One E1/PRI port at 75-ohm unbalanced
- Two E1/PRI ports at 75-ohm unbalanced
- One E1/PRI port at 120-ohm balanced
- Two E1/PRI ports at 120-ohm balanced
- One channelized E1 75-ohm unbalanced or 120-ohm balanced
- One T1/PRI port
- Two T1/PRI ports

Figure 2-26
The MIP card configurations.



These configurations allow the MIP to provide varied answers to connectivity requirements. The dual-port MIP can act as a dial-on-demand (DDR) ISDN PRI for high-volume locations or be configured through software, enabling one port to act as an ISDN PRI line while the other operates as a multichannel multiplexer feeding remote locations.


7x00 Packet OC-3 Interface Processor (POSIP)

The POSIP board, shown in Figure 2-27, complies with RFC 1619, "PPP over SONET/SDH" and RFC 1662, "PPP in HDLC-like Framing." Using these standards, the POSIP encapsulates packet data using Point-to-Point Protocol (PPP). The PPP is then mapped into an STS-3c/STM-1 frame, reducing the transport overhead by approximately 50 percent, as compared to using ATM adaptation Layer 5 (AAL5) and line card control (LCC) Sub-network Access Protocol (SNAP) encapsulations over SONET OC-3 media.

The POSIP interface supports one 155 Mbps port using either single-mode or multimode optical-fiber on Cisco 7000 and 7500 series routers. The Cisco 7000 must have the RSP7000 system processor installed to support the POSIP board. The POSIP has support for the following features:

- SONET/SDH-compliant interface, SONET/STS-3c, and SDH/STM-1 framing and signaling overhead
- Full-duplex operation at OC-3 155 Mbps

Figure 2-27
The POSIP card.



- Intermediate reach optical interface with single-mode fiber
- Optical interface with multimode fiber
- OIR

The POSIP board connects the OC-3 optical-fiber network to the CxBus on the 7000 series or the CyBus on the 7500 series routers. The POSIP installs on any available interface processor slot. The POSIP board can be configured with 16 or 32 MB of DRAM and 1 or 2 MB of SRAM. The memory requirements can be upgraded at a later date.

7x00 Service Provider MIP (SMIP)


Internet Service Providers (ISPs) require speed in delivering packets between the end user community and the Internet. The SMIP functions similarly to the MIP. However, the SMIP does not support multiprotocol routing. Using Cisco IOS Release 10.2(6) or later is required to support the following SMIP functions:

- IP routing with PPP or High-Level Data Link Control (HDLC)
- ISDN PRI connectivity

The SMIP, shown in Figure 2-28, supports three different types of configurations. These are

- Two T1 ports
- Two E1 ports with 75-ohm
- Two E1 ports with 120-ohm

Figure 2-28
The SMIP card.




Note that the SMIP is only optioned with two ports. One port can be used to channelize Nx64 or Nx56, supporting 24 channels on a T1 or 30 channels on an E1. Each channel is configured as its own serial interface. The second port can be used as an ISDN PRI port for ISDN BRI dial connections to the router.

7x00 Standard Serial Interface Processor (SSIP)

The SSIP is only optioned with eight high-speed serial ports. The total aggregate bandwidth supported by the SSIP is 8 Mbps. The dual-port port adapters used on the SSIP are compatible with the FSIP. They are not interchangeable with the VIP2 or 7200 series port adapters. Each port diagrammed in Figure 2-29, when using Cisco IOS Release 10.3(6) or later, supports up to T1 or E1 speeds when using IP routing encapsulated in PPP or HDLC. If multiprotocol routing is required, the serial port uses PPP or HDLC encapsulation with speeds at 64 Kbps or less.

Figure 2-29
The SSIP card.



7x00 Token Ring Interface Processor (TRIP)

The Token Ring Interface Processor (TRIP) connects the Cisco CxBus or CyBus to a Token Ring network at 4 or 16 Mbps. Each port is connected to a Token Ring multistation access unit (MAU) using a DB-9 connector. The TRIP is configurable with either two or four Token Ring ports. Figure 2-30 illustrates the TRIP board.

7x00 Versatile Interface Processor 2 (VIP2)

The VIP2, shown in Figure 2-31, is a new generation interface processor board with a high-speed RISC MIPS 4700 processor with an internal speed of 100 MHz and a system bus interface speed of 50 MHz. This CPU enables the VIP2 to process all functions on the VIP2, rather than requesting functions from the RSP system processor. This function is available with Cisco IOS Release 11.1(472) or later, enabling the VIP2 to run the Cisco IOS kernel directly on its own CPU. The 7000 and 7010 series routers must have the RSP7000 and 7000CI system boards installed in order to use the VIP2 features.

Figure 2-30
The TRIP card.





Figure 2-31
The VIP2 card.



The VIP2 is comprised of a motherboard and up to two port adapters or service adapters. Any combination of port or service adapters can be installed on the VIP2 in support of LAN and WAN interfaces and services. Appendix E details the VIP2 models of VIP2 required in support of various port adapter and service adapter configurations.

Cisco 7x00 Series Port and Service Adapters

The port and service adapters for the 7x00 series routers are compatible between the VIP2 and the 7200 series router. The 7000 and 7010 series routers must have the RSP7000 and 7000CI system boards installed prior to using the VIP2 board supporting the port adapter and service adapters.

The following media and interface types are supported on the entire 7x00 series product line:

- ATM
- 100VG-AnyLAN
- Ethernet 10BaseT
- 10BaseFL
- Fast Ethernet 100BaseTX
- 100BaseFX
- Token Ring
- Fiber Distributed Data Interface (FDDI)
- High-Speed Serial Interface (HSSI)
- Synchronous serial media
- Channelized T1/ISDN PRI

The Cisco 7200 series supports all of the above media and interface types along with support for ATM-Circuit Emulation Services (ATM-CES) and ISDN PRI and BRI connections.

7x00 ATM OC-3


The ATM OC-3 comes in two models, as shown in Figure 2-32. The port adapter uses a single-port SC duplex connector to the OC-3c ATM network. It is supported on the full 7x00 series line when used with Cisco IOS Release 11.1(9)CA. The fiber run from the router to the switch may be up to 15 km in length.

7x00 ATM-Circuit Emulation Services (ATM-CES)


The ATM-CES is supported only on the 7200 series routers. It supports four T1 CES interfaces and a single ATM trunk for servicing data, voice, and video traffic over an ATM WAN using Cisco IOS Release 11.1(11)CA or later. As shown in Figure 2-33, the ATM-CES can support either structured Nx64 Kbps or unstructured 1.544 Mbps circuits. The ATM-CES is optioned with either an OC-3 (155 Mbps) single-mode intermediate reach ATM trunk interface or a DS-3 (45 Mbps) ATM trunk interface.

Figure 2-32

The ATM-OC-3 card.

**Figure 2-33**

The ATM-CES card.



7x00 100VG-AnyLAN

The 100VG-AnyLAN standard was developed and published by Hewlett-Packard (HP). Its intention is to provide voice, video, and data transport over 100 Mbps using Ethernet. The 100VG-AnyLAN port adapter uses a single interface port supporting the IEEE 802.12 specification of running

802.3 Ethernet packets at 100 Mbps over Category 3 or Category 5 UTP cable with RJ-45 terminations. The 100VG-AnyLAN port adapter operates at 120 Mbps using the 5B/6B coding scheme to provide the 100 Mbps data rate at half-duplex. Figure 2-34 depicts the 100VG-AnyLAN port adapter.

7x00 ISDN Basic Rate Interface (BRI)

The ISDN BRI port adapter is available only on the 7200 series router. Using an Network Termination 1 (NT1) device, the 7200 ISDN BRI port adapter connects using either one or both of the two B channels (64 Kbps) in full-duplex mode, observing an aggregate rate of 128 Kbps. The single D channel on the BRI is also available at a full-duplex data rate of 16 Kbps. Figure 2-35 illustrates the two models available for the 7200 series router. The port adapters are available in either four or eight ISDN BRI ports. The four-port ISDN BRI port adapter uses a connect switch on a U interface, while the eight-port ISDN BRI port adapter uses an S/T interface to the NT1 device.

Figure 2-34
The 100VG-AnyLAN card.





Figure 2-35
The ISDN BRI card.



7x00 Channelized T1/E1 ISDN PRI

The channelized port adapters from Cisco support T1 (1.544 Mbps) and E1 (2.048 Mbps) line speeds with the capability to connect using ISDN PRI standards. Each port adapter is available with one or two interfaces. The channelized E1/ISDN PRI port adapter is available with unbalanced 75-ohm or balanced 120-ohm connections. Figure 2-36 illustrates the channelized T1/E1 ISDN PRI port adapter.

7x00 Ethernet 10BaseT

The IEEE 802.3 Ethernet 10BaseT standard is supported using four or eight interfaces. Each interface runs at a wire speed of 10 Mbps, thereby providing an aggregate bandwidth of 40 Mbps for the four port and 80 Mbps for the eight port. The Ethernet 10BaseT port adapter, depicted in Figure 2-37, is available on the entire Cisco 7x00 router platform.

Figure 2-36
The Channelized
T1/E1 ISDN PRI card.





Figure 2-37
The Ethernet 10BaseT four-port and eight-port card.



7x00 Ethernet 10BaseFL

Support for 10-Mbps Ethernet over fiber-optic media is provided by using the 10BaseFL port adapter. The port adapter has up to five interfaces using the IEEE 802.3 Ethernet 10BaseFL standard running at 10 Mbps, each in half-duplex mode with an aggregate bandwidth rate of 50 Mbps. The interfaces, as shown in Figure 2-38, use a pair of multimode S/T receptacles, one for receive (RX) and one for transmit (TX), with both at wire speed. The Ethernet 10BaseFL is supported across the Cisco 7x00 router platform.

7x00 Fast Ethernet

The Cisco Fast Ethernet port adapters support full- and half-duplex operations at 100 Mbps. This port adapter is available on all the Cisco 7x00 router platforms and comes in two models.

In support of twisted pair media, the Fast Ethernet port adapter provides a single 100BaseTX port for connection to Category 5 UTP media using an RJ-45 connection. The 100BaseTX port adapter, shown in Figure 2-39, can also connect to Category 3, 4, and 5 UTP or STP for 100BaseT4 media using the MII interface. Additionally, the 100BaseTX Fast Ethernet model can connect to multimode fiber for 100BaseFX media using the MII interface through external transceivers.

Figure 2-38
The Ethernet 10BaseFL card.





Figure 2-39

The 100BaseTX and 100BaseFX Fast Ethernet card.



Connectivity to fiber-optic media is also available using the 100BaseFX Fast Ethernet port adapter. The 100BaseFX port adapter, shown in Figure 2-39, connects to fiber-optic media in one of two ways. The 100BaseFX can use SC fiber-optic connectors or use external transceivers to multimode fiber through the MII interface. Additionally, the 100BaseFX Fast Ethernet port adapter allows connectivity to 100BaseT4 networks through the MII interface over Category 3, 4, and 5 UTP or STP media.


7x00 Synchronous Serial

The synchronous serial port adapter comes with four interfaces. Each interface must be alike and supports the following electric standards:

- EIA/TIA-232
- EIA/TIA-449
- EIA-530 X.21
- V.35

The interfaces support either DCE or DTE terminations, depending on the type of cable connected to the interface. The synchronous serial port adapter depicted in Figure 2-40 is available on the Cisco 7500, 7000, and 7200 series routers.

Figure 2-40
The synchronous serial card.



7x00 Single Port Molex 200-pin Receptacle

The Molex 200-pin receptacle supports a wide variety of synchronous serial interfaces. Each Molex receptacle interface provides up to eight synchronous serial interfaces using a special cable designed for supporting the desired electrical interface specification. The Molex runs full-duplex mode, supporting either 1.544 Mbps (T1) or 2.048 Mbps (E1) speeds for V.35 and X.21 interfaces. Support for EIA/TIA-232 interfaces allows up to eight ports operating in full-duplex mode at 64 Kbps. Figure 2-41 illustrates the 200-pin Molex receptacle. These port adapters are available on the 7x00 family of routers.

7x00 Synchronous Serial E1-G.703/G.704

The E1-G.703/G.704 serial interface is an International Telecommunication Union Telecommunication (ITU-T) standard for serial line speeds of 2,048 Mbps on E1 lease lines. The port adapter supports up to four synchronous serial interfaces for framed and unframed service. The interfaces are ordered with eight unbalanced 75-ohm or balanced 120-ohm. Figure 2-42 diagrams the Synchronous Serial E1-G.703/G.704 port adapter.

7x00 Token Ring

The Token Ring port adapter provides up to four IEEE 802.5 Token Ring interfaces at either four or 16 Mbps. The port adapter is available on the

Figure 2-41
The 200-pin Molex receptacle.





Figure 2-42

The Synchronous Serial E1-G.703/G.704 port adapter.



7x00 family of routers and comes in two models. A half-duplex and full-duplex model. The full-duplex model realizes an aggregate speed of 32 Mbps. Figure 2-43 illustrates the Token Ring port adapter.

7x00 FDDI

The FDDI port adapter comes in two flavors: half-duplex and full-duplex. Each of these flavors is available with two multimode or single-mode interfaces at a maximum bandwidth of 100 Mbps per port. Each port adapter supports the optical bypass switching capability. Figure 2-44 diagrams the single and multimode FDDI port adapters. The full-duplex option enables the FDDI port adapter to realize and aggregate speed of 200 Mbps per port. The FDDI port adapters are available for the entire Cisco 7x00 family of routers.

Figure 2-43

Token Ring port adapter.





Figure 2-44

The single and multimode FDDI port adapters.



7x00 HSSI

HSSI port adapters are configurable with either one or two HSSI interfaces. Each interface uses the EIA/TIA 612/613 high-speed standard to provide T3 (45 Mbps), E3 (34 Mbps), and SONET STS-1 (51.82 Mbps) data rates. Figure 2-45 illustrates the HSSI port adapter, which is available on all Cisco 7x00 routers.

7x00 Compression Service Adapter

Bandwidth for many installations is a valuable asset. Compressing data prior to transmission enables routers to transmit more information than would be allowed without compression. The compression service adapters off-load compression and decompression functions from the host processor for inbound and outbound traffic over channelized E1/ISDN PRI, channelized T1/ISDN PRI, BRI ISDN, and synchronous serial port adapters. Figure 2-46 diagrams the two models for the compression service adapters.

The first model has 786 KB of memory, enabling it to handle compression/decompression up to 64 WAN links. The second model is configured with 3

Figure 2-45
The HSSI port adapter.





Figure 2-46
The compression service adapters.





MB of memory in support of 256 WAN links. Both models of the compression service adapter are available on the entire Cisco 7x00 family of routers.

4000 Series

The Cisco 4x00 router platform is based on the use of network processor modules (NPM). Using the NPMs, a 4x00 router can combine many different types of interface connections in support of various networking requirements. The 4x00 series router platform is available in three models. Each model looks identical, as depicted in Figure 2-47, with different interface support and processing power. The models 4000-M, 4500-M, and 4700-M can mix and match the NPMs using the three available slots. The low-end 4000-M model supports the following NPMs:

- Ethernet
- Token Ring
- FDDI
- Serial
- ISDN BRI
- Channelized E1/T1 ISDN PRI

Figure 2-47
The Cisco 4000 series router.



The higher-end 4500-M and 4700-M routers support the following network interfaces in any combination using the three available slots:

- Ethernet
- Token Ring
- FDDI
- HSSI
- High-density serial
- ISDN BRI
- Channelized E1/T1 ISDN PRI
- ATM OC-3c
- ATM DS-3
- ATM E3

The NPMs available for each router platform come in various port configurations. Although some have multiple ports, the 4000 series platform supports full wire speed on each port. Each NPM has the following port configurations:

- One-, two-, or six-port Ethernet
- One-port Fast Ethernet
- One- or two-port Token Ring
- One-port multimode FDDI (both single [SAS] and dual attachment station [DAS])
- One-port single-mode FDDI (DAS)
- Two- or four-port synchronous serial
- Two-port high-speed serial and 16-port low-speed serial
- Four- or eight-port ISDN BRI
- One-port channelized T1/ISDN PRI
- One-port channelized E1/ISDN PRI (balanced or unbalanced)
- Four-port serial G.703 and G.704 (balanced or unbalanced)
- One-port HSSI
- One-port ATM (single-mode or multimode) OC-3c
- One-port ATM DS-3
- One-port ATM E3

Due to the processing of the high-speed NPMs, a maximum of two high-speed interfaces is available on the Cisco 4500-M and Cisco 4700-M platforms. This means only two of the following NPMs can be installed and operable using the Fast Ethernet, FDDI, ATM-OC3, or DS-3 NPMs. The exception to this is that there can only be one ATM-OC3 NPM configured and operable on the 4500-M or 4700-M routers. Therefore, combinations with the ATM-OC3 NPM are

- One Fast Ethernet
- One FDDI
- One ATM-DS3 or E3
- One HSSL

The 4500 or 4700 routers, however, can be configured with

- Two Fast Ethernet
- Two FDDI, two HSSI
- One Fast Ethernet and one FDDI
- One Fast Ethernet and one HSSI
- One FDDI and one HSSI.

In these types of configurations, the remaining slot can be used by the other NPMs, as noted. For complete detail of NPM configurations and combinations, see Appendix F.

4000 Usage

The 4000 series routers were initially developed as access routers in the Cisco routing architecture. The 4700-M router using the high-speed NPMs can perform the duties of a distribution router as well as an access router.

4000 Processors

The processors vary on each platform. The 4000-M series uses a Motorola 40 MHz 68030 processor, while the 4500-M and the 4700-M uses an IDT Orion RISC processor. The Cisco 4500-M router uses a 100 MHz IDT Orion RISC processor, while the high-end 4700-M platform uses a 133 MHz IDT Orion RISC processor.

4000 Memory

Each 4000 series router comes standard with 128 KB of NVRAM, which is used to store and recall the router configuration. Main memory on the router is used for executing the Cisco IOS and process routing tables. Shared memory is used to move packets between interfaces, and flash memory is used to store router configurations and Cisco IOS code. Since the 4000 series is actually designed for the access layer of the Cisco routing architecture, it comes with low base memory.

The 4000-M platform comes with a base of 4 MB of flash memory, expandable to either 8 MB or 16 MB. Main memory on the 4000-M starts with 8 MB and can be expanded to 16 or 32 MB of memory. Earlier models of the the 4000-M shipped with 1 MB of shared memory, while the newer models are shipped with 4 MB of shared memory. If the 4000-M being used is an earlier model, the shared memory must be upgraded to a minimum of 4 MB to support FDDI or have more than five physical or virtual interfaces defined. Shared memory is expandable to 16 MB.

The flash memory support on the 4500-M platform is the same as that found on the 4000-M router. Main memory comes standard at 16 MB and with an upgrade to 32 MB of main memory. The 4500-M router comes standard with 4 MB of shared memory with the option to expand to 8 or 16 MB.

The 4700-M platform also comes standard with 4 MB of flash memory, with upgrades to either 8 or 16 MB. Being the high end of the 4000 series platform, the 4700-M comes standard with 16 MB of main memory with expansion to either 32 or 64 MB of memory to handle large routing tables. Shared memory on the 4700-M is the same as that found on the 4500-M router.

3600 Series

The 3600 series router is one of the newest modular platforms from Cisco. This router comes in two models: the 3640 and the 3620. The 3600 series provides for increased dial-up port density with newer WAN technologies like ATM. One special feature available on the 3600 series is the capability for the operator and auxiliary consoles to connect to a local or remote PC at 115.2 Kbps. This allows support for Xmodem or Ymodem protocols in order to load the router IOS software directly through these ports, versus having to have a network connection.

The 3640 has more port capacity than the 3620, as shown in Figure 2-48. The 3640 is available with four network module slots, while the 3620 has two network module slots available. The module slots are used to connect external media to the bus backplane of the router with network module interface cards that mix LAN and WAN media types, along with asynchronous and synchronous serial connections and support for ISDN PRI and BRI interfaces.

In support of ISDN PRI connectivity, the 3640 installed with a mixed media module and three two-port ISDN PRIN network module interfaces can connect up to 138 T1 or 180 (E1) B channels. This establishes the 3640 as a cost-effective solution for corporate telecommuting. Using three eight-port ISDN BRI network interface modules, the Cisco 3640 can connect up to 48 B channels with local LAN and WAN routing capability.

The port density on the network interface cards enables the 3640 to support up to 24 asynchronous or synchronous serial interfaces for multiple 56 Kbps connections. The 3600 series routers support the following network interfaces:

- One- and four-port Ethernet network modules
- One-port Fast Ethernet network module
- One-port Ethernet and one-port Token Ring network modules
- Four- and eight-port Asynchronous/Synchronous network modules
- Four-port serial network module
- ISDN BRI (ST and U interfaces)
- Channelized T1/ISDN PRI (with and without CSU)
- Channelized E1/ISDN PRI (balanced and unbalanced)


The 3600 series of Cisco routers requires Cisco IOS software Release 11.1(7)AA and later, or Release 11.2(5)P and later.

The network modules are the cards that slide into the slots of the 3600 series routers, as shown in Figure 2-49. The network modules themselves provide various interfaces for connecting external networks to the router bus backplane. Of these network modules, one of the more versatile is the mixed-media network module.


The mixed-media network module supports up to two fixed LAN interfaces and two user-installable WAN interfaces. The LAN interfaces are a part of the network module itself and cannot be removed. The LAN interface support, as illustrated in Figure 2-50, is one of the following:

- One Ethernet port
- Two Ethernet ports
- One Ethernet and one Token Ring port


Figure 2-48
Cisco 3600 series
routers.




H7221



H6551




H7336



H7238

Figure 2-49
Network Module
cards for the 3600
series routers.




The Ethernet connections support both 10BaseT and AUI interfaces at 10 Mbps. The Token Ring port is either four or 16 Mbps, using either STP or UTP wiring. The WAN expansion slots on the mixed-media network module support the following WAN interface cards:

- One-port ISDN BRI WAN interface card
- One-port ISDN BRI with NT1 WAN interface card
- One-port serial WAN interface card
- One-port four-wire 56 Kbps DSU/CSU WAN interface card

Each of the WAN network interface cards is shown in Figure 2-51. The Cisco 1600 series routers also support the

- Cisco 3600 ISDN BRI
- ISDN with NT1
- Serial interface cards.


Figure 2-50
LAN interface support
on 3600 series
routes.



The 3600 series router requires Cisco IOS Release 11.2(4)XA or 11.2(5)P or later to properly operate the

- WAN interface cards ISDN BRI
- ISDN with NT1
- One-port four-wire 56 Kbps DSU/CSU interface cards.

Figure 2-51
WAN interface cards
for the 3600 series
router.



The network modules supporting channelized T1/E1 and ISDN PRI lines are available with a built-in CSU with one or two ports. Figure 2-52 illustrates the various channelized T1/ISDN-PRI and E1/ISDN-PRI network modules available for the 3600 series routers.

Using a T1/ISDN-PRI CSU, the network module connects directly to the provider's network connection. Without the internal CSU, the T1/ISDN PRI network module connects to an external CSU, which then connects to the provider's network connection. The T1 module channelizes the T1 up to 24 virtual channels per T1 port. The E1/ISDN PRI network module provides one or two E1 ports at 2.048 Mbps second in full-duplex transmission. They are configured as either balanced or unbalanced and provide up to 30 virtual channels per E1 port. If the T1/E1 modules are configured for using ISDN PRI, they are not compatible with the four or eight-port ISDN BRI modules. However, when used as a "multiplexer," the ISDN BRI modules are compatible.

The ISDN BRI network modules have four different models. They use four or eight ISDN BRI ports, along with S/T or on-board NT1 service, and each port defines the four different model types, as shown in Figure 2-53. The ISDN BRI network modules use local SRAM for buffer descriptor, input queues, and configuration storage to increase performance. The performance of the ISDN BRI eight-port model is 5,760 packets per second (pps), running full-duplex continuous data of 144 Kbps using 50-byte packets. The aggregate full-duplex rate of the eight-port ISDN BRI network module is 2.3

Figure 2-52
The various channelized T1/ISDN-PRI and E1/ISDN-PRI network modules available for the 3600 series routers.

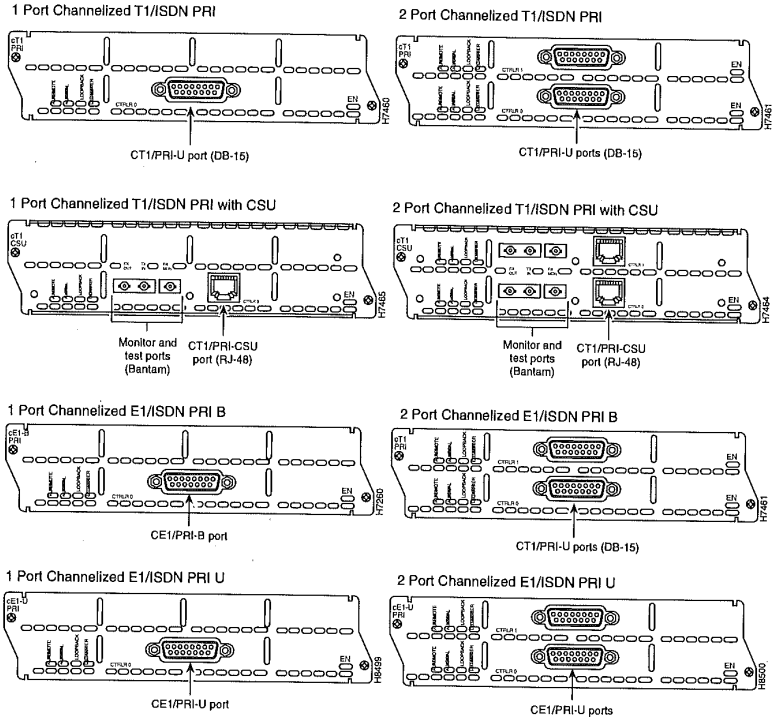



Figure 2-53
The four ISDN BRI network modules for the 3600 series router.




Mbps. The ISDN BRI network modules include features to query the network module, SNMP traps for monitoring the network module, manageability with Ciscoworks or CiscoView, and support for the ISDN MIB standard.


For more traditional low-speed network connections, the four- and eight-port Asynchronous/Synchronous network modules are available. Figure 2-54 illustrates the two module formats. These network modules support 128-Kbps synchronous connections or 115.2-Kbps asynchronous connections per port. The ports use the DB-60 interface standard for connecting to the router.

In support of Ethernet, the 3600 series network modules are available with one- and four-port Ethernet connections. As shown in Figure 2-55, the one-port Ethernet network module comes with one AUI DB-15 and one 10BaseT RJ-45 interface connection. Only one of these ports can be active at any time for this network module. The four-port Ethernet adds to the one-port Ethernet network module format with three 10BaseT RJ-45 connections on the left side of the network module. The restriction of either the

Figure 2-54
The
Asynchronous/Synchronous network
modules.




4-port Asynchronous/Synchronous



8-port Asynchronous/Synchronous

Figure 2-55

Ethernet network module support on the 3600 series router.



AUI or RJ-45 port being active on the right side of the four-port Ethernet module still holds true. Cisco IOS Release 11.2(4)XA and 11.2(5)P or later are required for operation.

The advancement of Ethernet has dictated that the network modules keep with the new Ethernet standards. Currently, the 3600 series routers support a one-port Fast Ethernet network module using an RJ-45 connector or a 40-pin media-independent interface (MII). Again, there is a restriction that only one of these interfaces can be active at any given time. The RJ-45 connects two pairs of Category 5 UTP wiring using the 100BaseTX standard. Using the MII, an external transceiver is required to connect to a multimode optical fiber using 100BaseFX standard or it can use the 100BaseT4 standard over four pairs of Category 3, 4, or 5 UTP or STP wiring. Figure 2-56 diagrams the one-port Fast Ethernet network module for the 3600 series router, which requires Cisco IOS Release 11.2(6)P or higher for operation.

The Cisco 3640 supports a maximum of two one-port Fast Ethernet network modules with no other network modules installed. If using the one-port Fast Ethernet with a four-port Ethernet network module, the 3640 router can be configured for a maximum of one Fast Ethernet and two four-port Ethernet network modules, along with other network modules.

Using the high-density DB-60 interface standard, the four-port serial network module can support various data rates. Figure 2-57 illustrates the four-port serial network module. If only port 0 is used, then the interface can realize a data rate of 8 Mbps. Using ports 0 and 2, the data rate is halved to 4 Mbps per port, and the data rate is halved again to maximum of 2 Mbps per port when using all four ports.

3600 Usage

The 3600 series routers are designed for the access layer of the Cisco router architecture. The 3640 is ideal for use by ISPs to have many points of presence (POPs) or for telecommuting to a corporate environment. The 3620 provides for small-office connectivity and local LAN and WAN connections using mixed media network modules.

Figure 2-56
Fast Ethernet network module for the 3600 series router.





Figure 2-57
Four-port serial interface network module.



3600 Processors and Memory

The two models of the Cisco 3600 series use different processors. The Cisco 3640 uses the 100-MHz IDT R4700 RISC processor, and the Cisco 3620 uses the 80-MHz IDT R4600 RISC processor. The 3600 series uses a single DRAM pool, which is partitioned main and shared memory areas. This partitioning of DRAM makes memory calculation difficult when configuring the 3600 router platforms. Appendix F identifies some guidelines on how to configure the proper amount of DRAM for the 3600 routers.

The 3600 series also uses flash memory. Both the DRAM and flash utilize the SIMM chips for memory, allowing field upgrades and replacements. The standard flash memory is 4 MB, but the flash memory can be upgraded to a maximum of 48 MB for both the 3620 and 3640 routers. Each router comes with a base of 16 MB of DRAM, which is expandable on the 3620 to 64 MB and on the 3640 to 128 MB. In addition to on-board flash memory, the 3600 series has two slots of PCMCIA available in support of 4 MB to 128 MB of flash using two 64-MB PCMCIA flash cards.

2600 Series

The Cisco 2600 series router platform extends the modular format of the 3600 series into the smaller remote branch office. The modularity of the 2600 series enables these small offices to deploy voice/fax/video along with data in a single versatile network appliance. The Cisco 2600 shares many of the same network module interfaces with the 3600 and 1600 router platforms. The 2600 series supports one network module slot, two WAN Interface Card slots, and a new interface slot dubbed Advanced Integration Module (AIM). Cisco maximizes uptime on the 2600 series through the use of an external Redundant Power Supply (RPS) and Cisco IOS dial-on-demand routing features for the restoration of both data and voice connections automatically, should the primary link failure occur.

The 2600 series comes in two flavors: a single Ethernet (2610) or a dual Ethernet interface (2611). The WAN interface card slots support the following:

- Serial
- ISDN BRI
- Built-in CSU/DSU functions

The network modules add needed support for

- Multiservice voice/data/fax integration
- Departmental dial concentration
- High-density serial concentration

The AIM slot supports added features for optimization through hardware-assisted data compression and encryption.

An auxiliary port with the capability for use as a 115-Kbps DDR interface for WAN backup connectivity is standard on both the 2610 and 2611 models. Figure 2-58 shows the rear panel of the 2600 models.


The Cisco 2600 shares many of the data network modules with the 3600 series routers. These shared data network modules are

- 16-port high-density async network module
- 32-port high-density async network module
- Four-port low-speed (128 Kbps max) async/sync serial network module
- 8-port low-speed (128 Kbps max) async/sync serial network module

The following voice/fax network modules and interface cards are shared with the 2600 and 3600 series router:

- One-slot voice/fax network module
- Two-slot voice/fax network module
- Two-port FXS voice/fax Interface card
- Two-port FXO voice/fax interface card
- Two-port E/M voice/fax interface card

Figure 2-58
The rear panels of
the 2600 series
router.



The 2600 series also shares WAN Interface Cards (WICs) with the 1600 and 3600 series routers. These cards are the following:

- One-port serial WIC
- One-port four-wire 56 Kbps DSU/CSU
- One-port ISDN BRI
- One-port ISDN BRI with NT1

WICs unique to the 2600 series support the following configurations:

- Two-port serial WIC for Cisco 2602
- Two-port async/sync serial WIC for Cisco 2602

2600 Usage

Based on its size and purpose, we can see that the 2600 series falls into the access layer of the Cisco-layered network topology. Multiservices have become quite desirable for reducing communications network infrastructure cost while at the same time enhancing application functionality. Using the QoS features built into the Cisco IOS software, small branch offices can participate in voice-enabled desktop applications and desktop video.

Using the modular features, the 2600 can serve as a dial services concentrator for remote office and remote user access by supporting up to 36 high-speed asynchronous ports using

- PPP
- SLIP
- ARA
- Xremote protocols.

This enables casual connection for these remote locations to the corporate WAN through the WAN interface cards available on the two 2600 models.

The various WAN modules and WIC slot options enable the 2600 series routers to be a serial device concentrator. Through the power of the Cisco IOS and optional support of up to 12 synchronous serial interfaces, the 2600 protects legacy system investment for SDLC, bisynch, and asynch devices. Ideally, this capability in combination with the Ethernet LAN interfaces and integrated CSU/DSU and ISDN BRI WAN interface cards allows a network designer to provide a solution for connecting retail, financial, and sales branch offices.

2600 Processor

The 2600 series router has a Motorola MPC860 40-MHz CPU with a 20-MHz internal bus clock.

2600 Memory

The system memory (DRAM) comes in two DIMM slots. The default memory size is 16 MB with a possible expansion of a total of 64 MB. Flash memory is incorporated on the processor board using a single SIMM slot supporting a default of 4 MB with expansion to 16 MB. The DRAM on the 2600 uses pooled DRAM memory, which is partitioned between processor and packet memory areas. The default 16 MB of DRAM is partitioned into 12 MB for processor and 4 MB for packet memory.

Regarding Cisco IOS Release 11.3(2)XA, 11.3(3)T, and higher, the Cisco IOS can be loaded into the router using the LAN interface and TFTP or by using the auxiliary or console port with the Ymodem or Xmodem protocols. This is valuable for remote dial-up restoration of a damaged IOS or for updating the stored configuration file.

2500 Series

The 2500 series router platform from Cisco provides specific access layer functions for small offices or small business. The 2500 series comes in many different solution formats:

- Single LAN routers
- Mission-specific routers
- Router/hub combinations
- Dual LAN routers
- Modular routers

Additionally, the 2500 series comes in an access server for supporting remote dial-up access to enterprise networks. The Cisco access servers are not discussed in this text. The console and auxiliary ports on the 2500 series use RJ-45 connectors. Any 2500 series model ordered comes with a cable kit to connect an RJ-45-to-RJ-45 using a roll-over console cable, an RJ-45-to-DB-25 male DCE adapter, an RJ-45-to-DB-25 female DTE adapter, and an RJ-45-to-DB-9 female DTE adapter for connecting PCs or

modems to these ports. The low-speed serial asynch/synchronous ports on all the models support asynchronous connections up to 115.2 Kbps and synchronous connections up to 2 Mbps.

The single LAN routers come in eight models. Each model has a different combination of non-upgradeable or non-field modifiable interfaces. The 2501, shown in Figure 2-59, provides a single Ethernet 10-Mbps port and two synchronous serial interfaces. The Ethernet uses a DB-9 AUI port, which may require an external transceiver to connect to an RJ-45 LAN hub interface. The two serial ports use DB-60 connectors and all data rates up to 2 Mbps.

The 2502 router pictured in Figure 2-59 has a Token Ring LAN interface, instead of an Ethernet AUI port. The Token Ring interface uses a DB-9 connection, which may require a converter to an RJ-45 connector for connecting to a LAN hub. The Token Ring interface is configurable as 4 or 16 Mbps data rates.

The addition of a single ISDN BRI port on the router is shown on the 2503 and 2504 routers in Figure 2-59. Note that the 2503/2504 is the same as the 2501/2502 with the exception of the ISDN BRI ports. The ISDN BRI ports have an internal ISDN Terminal Adapter. These ports must connect to an ISDN NT1 device for switched ISDN connectivity.

Support for low-speed asynch/synchronous serial lines is provided by the 2520/2521 platforms pictured in Figure 2-60. There are two low-speed connections with asynchronous data rates up to 115.2 Kbps and synchronous data rates up to 128 Kbps. Additionally, the LAN ports for Ethernet and Token Ring are also provided with an RJ-45 connection interface. Only one LAN interface is allowed to be configured and operative at any one time.

The 2520/2521 also provides a single ISDN BRI port. The 2520 Ethernet AUI or 10BaseT RJ-45/UTP adapter supports 10 Mbps and the 2521 Token Ring DB-9 or RJ-45/UTP adapter supports 4 or 16 Mbps data rates.

The last two models, pictured in Figure 2-60, in the single LAN category of the 2500 series routers provide for up to eight low-speed asynch/synchronous and two high-speed communications interfaces, a single ISDN BRI, and a single LAN interface. The 2522 provides for Ethernet at 10 Mbps using an AUI or a 10BaseT RJ-45 connection. The 2523 model supports the 4 or 16 Mbps Token Ring speeds using either the DB-9 or RJ-45 UTP ports.

Mission-specific entry-level routers in the 2500 series come in 12 unique offerings. The mission-specific router models are configured with less memory than the single LAN models and execute IOS software images specifically designed for the CFRAD (CF), LAN FRAD (LF), and ISDN requirements. The special IOS images disable/enable unused ports through

Figure 2-59
Cisco
2501/2502/2503/
2504 series routers.





Figure 2-60
The
2520/2521/2522/2523
models of the
Cisco 2500 series
routers.



software. These mission-specific routers give the single LAN router platform the capability to act as frame-relay access devices for connecting the location to frame-relay networks without having to connect through a separate frame-relay access piece of equipment. The CF models allow the router to also act as a frame-relay switch for delivering information through frame-relay networks. These models, however, are upgradeable to full functionality through full-function IOS software and added memory.

The mission-specific routers are the exact models of the single LAN routers, though the software has limited functionality. The 2501CF/2502CF routers have their respective LAN ports disabled by the IOS software and only allow configuration of the two high-speed serial interfaces. The 2501LF/2502LF have their LAN ports enabled along with the capability to send LAN traffic through frame-relay networks directly.

The 2503L/2504I provide for Ethernet and Token Ring LAN connectivity, respectively, through ISDN BRI connections. The high-speed serial connections available on the router are software-disabled.

The 2520CF, 2521CF, 2522CF, and 2523CF routers all have their ISDN BRI ports disabled and their respective LAN interfaces disabled as well. The low- and high-speed ports are enabled and functional.

The 2520LF, 2521LF, 2522LF, and 2523LF have all their LAN and WAN ports enabled; however, their ISDN BRI ports are disabled by the software.

For locations in which a single device must support both routing and LAN connectivity for workgroups and small offices, the Cisco 2500 series router/hub combination is available in six different formats. Each format supports only one LAN segment but has multiple ports available for connecting workstations or servers. The integrated hubs on these router platforms save the small business or small office equipment and software costs while providing a full LAN/WAN solution.

The 2505, 2507, 2516, and 2518 router/hub offerings, diagrammed in Figures 2-61 and 2-62, provide a single segment Ethernet LAN environment.

The 2505 supports up to eight Ethernet connections, the 2507 supports 16, the 2516 supports 14, and the 2518 supports 23 Ethernet LAN connections to the hub. The router card of the 2518 connects to port 24 of the Ethernet hub, allowing the 2518 to route LAN traffic over the WAN. The AUI port on the 2518 allows the 2518 to connect to an external Ethernet hub, expanding the reach of the LAN segment. Both the 2516 and the 2518 have the capability to expand to five hubs using Lanoptics hub expansion units. Each platform has two high-speed serial interfaces, and only the 2505/2507 do not provide for an ISDN BRI interfaces.

The 2517 and 2519, shown in Figure 2-63, support Token Ring LAN segments. The 2517 model allows for 11 Token Ring LAN connections to the

Figure 2-61

The 2505/2507 models of the Cisco 2500 series router.




Figure 2-62

The 2516/2518 models of the 2500 series router.





Figure 2-63

The 2517/2519 models of the 2500 series router.




hub, while the 2519 supports up to 23 Token Ring LAN segments to the hub. The hub interfaces can either be 4 or 16 Mbps, but all the ports must use the same data rate. The 2519 contains Token Ring ring-in/ring-out ports for cascading Token Ring hub equipment, thereby increasing the size of the Token Ring segment. Additionally, the ring ports 1 through 12 can be defined as a separate Token Ring segment from ports 13 through 24.

Both the 2517 and 2519 also have router cards with Token Ring RJ-45 connectors. The router cards attach to port 12 of the 2517 and port 24 of the 2519 routers. This enables the routers to transport LAN traffic over a WAN.

The 2517 allows a single port on the 11 available ports to connect to another hub using an RJ-45 cross-over cable expanding the Token Ring segment. On the 2519, the ring-in/ring-out ports allow for the expansion of the segment. An expansion unit is found on the top of the 2517 and 2519 to expand the hub to five hubs using Lanoptics supplied hubs. Both of these models have a single ISDN BRI port for switched backup use or bandwidth-on demand use in conjunction with two high-speed serial ports.

Small offices requiring more than one LAN are supported by the dual LAN router models. These are available in three different models. Figure 2-64 depicts the three dual LAN routers. All three models do not have ISDN BRI ports available. The 2513 supports one Ethernet 10 Mbps LAN segment and one Token Ring 4- or 16-Mbps LAN segment with two high-speed serial interfaces. The 2514 supports two Ethernet 10-Mbps LAN segments using AUI ports, and the 2515 supports two Token Ring LAN segments at 4 or 16 Mbps using DB-9 connectors.

Figure 2-64
The three dual LAN routers of the 2500 series router.




The modular routers in the 2500 series give the network engineer the ability to change and adapt the 2500 series routers, unlike the previous models mentioned. Two types of modular 2500 series routers exist. The two modular router models shown in Figure 2-65 differentiate themselves by the LAN support.

The 2524 connects Ethernet LANs, while the 2525 provides Token Ring connections. Both allow up to three WAN modules, configuring up to two synchronous serials and one ISDN. The modules are available in the following configurations:

- Two-wire, switched, 56-kbps DSU/CSU
- Four-wire, 56/64-kbps DSU/CSU
- Fractional T1/T1 DSU/CSU
- Five-in-one synchronous serial

Figure 2-65

The two 2500 series modular routers.



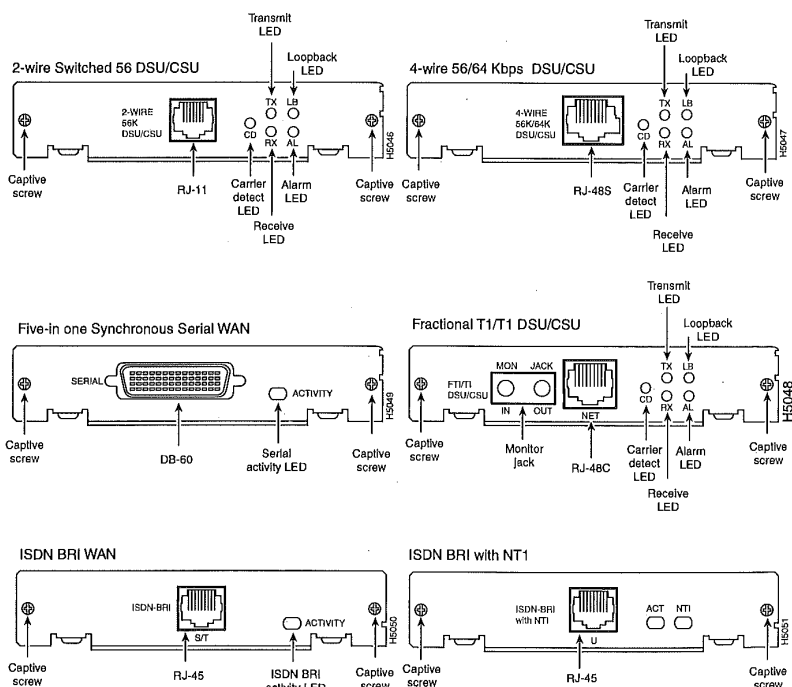
- ISDN BRI
- ISDN with integrated NT1 device

The three available slots, shown in Figure 2-65, on the 2524 and 2525 are used for the WAN interfaces. The WAN slot on the right of the unit is keyed to allow only the ISDN BRI interface cards to be installed. Likewise, the ISDN BRI cannot be installed in the first two WAN slots starting on the left of the router. The two-wire switched 56-Kbps DSU/CSU WAN module allows for 56-Kbps dial-up connections through the plain old telephone service (POTS) using an RJ-11 connection. The module connects directly from the RJ-11 port on the module to the RJ-11 port on the wall for connecting to the public telephone network.

The four-wire 56/64-Kbps DSU/CSU WAN module, shown in Figure 2-66, for the 2524 and 2525 router provides dedicated leased line synchronous serial connections up to 64 Kbps using and RJ-48S connector directly to the wall plate connecting the line to the communications network.

The fractional T1/T1 DSU/CSU WAN module, shown in Figure 2-66, uses an RJ-48C connector to the network. This module supporting a 1.544 Mbps

Figure 2-66
Network modules for the 2524/2525 routers.



line provides either Nx56 or Nx64 channels up to a total of 24 individual channels at each speed. Each channel is defined as if it were its own unique interface.

The ISDN modules pictured in Figure 2-66 provide ISDN BRI connectivity using RJ-45 S/T connections. The ISDN BRI supports two B channels and one D channel. The two B channels together allow for a switched connection of 128 Kbps. The ISDN BRI module contains its own terminal adapter and must be connected to an external NT1 device. The second ISDN BRI module has an integrated NT1 device and connects directly to the ISDN BRI port installed by the network provider.

The five-in-one synchronous serial WAN module, shown in Figure 2-66, enables the one interface to support the following electrical interface standards using the appropriate cables:

- EIA/TIA-232
- EIA/TIA-449
- V.35
- X.21
- EIA-530

The router side of the cable used has a DB-60 connector. The opposite end is headed with the appropriate interface required as specified by the line connection requirements.

2500 Usage


The 2500 series has many different uses and in some ways can provide both distributed and access layer functions. For example, a 2525 can connect a location to a frame-relay network with a 56-Kbps switched dial backup line to another 2525 at a different location. Meanwhile, a 2519 at a third site connects a Token Ring LAN to a corporate center using a 256-Kbps line to a multiplexer, which is attached to a 2424 with a Fractional T1 WAN module servicing all three remote sites. The WAN module also connects them to a core router in the larger corporate backbone.

2500 Processor and Memory

All the 2500 series router platforms use the Motorola 20 MHz 68030 processor. Each system comes with a minimum of 8 MB of flash memory.


the same ISDN line. Figure 2-67 illustrates the front of all the 1600 routers and the rear views of the four individual offerings. The expansion slot of the 1603 and 1604 is not available for a second ISDN port. However, the 1601 and 1602 can mix and match all the available WAN modules for the expansion slot.

Figure 2-67
The 1600 series routers.




Front of Cisco 1600 Series


Cisco 1601




Cisco 1602



Cisco 1603



Cisco 1604



Three WAN interface expansion modules are available with the 1600 series routers. Figure 2-68 diagrams their interface plates. The serial WAN interface expansion module provides EIA/TIA-232, V.35, X.21, EIA/TIA-499, and EIA-530 standard interfaces with support for 115.2-Kbps asynchronous and up to 2.048-Mbps synchronous connections. The proper cable must be installed to support the various interface requirements for successful operation. The ISDN BRI S/T supports two B channels and one D Channel for data only. The ISDN BRI U with a built-in NT1 allows connectivity to the a switched ISDN network without the use of an external NT1 device.


1600 Usage

The 1600 series routers are an ideal low-cost solution for small remote sales offices or telecommuters with a need for high-speed connectivity or casual connectivity to a single Ethernet LAN segment with IP/IPX or AppleTalk communication requirements. The 1600 series is the quintessential access layer router.

1600 Processor and Memory

The 1600 series uses the Motorola 68360 33 MHz processor. Each unit comes with a base of 4 MB of flash, which is expandable to 12 MB. Flash

Figure 2-68
The WAN serial interface expansion modules.



expansion can go from 4 to 6 MB or 4 to 8 MB or 4 to 12 MB. The DRAM comes with a base of 2 MB of memory, expandable to a maximum of 18 MB.

700M Family of Access Routers

The Cisco 700M family is an ISDN multiprotocol access router. The 700M family supports ISDN basic rate interface (BRI) of 56, 64, or 128 Kbps remote access connections. The Cisco 700M family of access routers comes in two series: the 760 and 770. The 760 series has one Ethernet 10-Mbps LAN interface and an ISDN BRI port. The 770 series includes a built-in four-port 10 Mbps Ethernet hub and a ISDN BRI, along with a call connect/disconnect switch on the format of the router to allow the user to manually connect or disconnect the ISDN BRI data linen connection.

The 760/770 series is broken further down into four models. Their features and functions are as follows:

■ 761M/771M

Shown in Figure 2-69, these models require an external NT1 device for connectivity. It is based on the Intel 25 MHz 386 processor and comes with 1.5 MB, which is expandable up to 2 MB over DRAM. The on-board NVRAM is 16 KB with 1 MB of flash memory. It can support up to 1,500 users and is available worldwide.

■ 762M/772M


Shown in Figure 2-70, these models include an internal NT1 device for connectivity. Additionally, these models have a second BRI port for external ISDN device connectivity or a second ISDN BRI line. It is based on the Intel 25 MHz 386 processor and comes with 1.5 MB, expandable to 2 MB over DRAM. The on-board NVRAM is 16 KB with 1 MB of flash memory. It can support up to 1,500 users and is available in North America only.

■ 765M/775M

Shown in Figure 2-71, these models require an external NT1 device for connectivity. It also includes two analog POTS RJ-11 ports for attaching phones, fax machines, and modems to share the ISDN BRI simultaneously with data. This model also supports provider supplemental services over ISDN such as call waiting, cancel call-waiting, call retrieve, call hold, three-way call conferencing, and call

Figure 2-69
The 761M/771M
routers.

Cisco761M



Cisco771M





Figure 2-70
The 762M/772M
routers.

Cisco 762M



Cisco 772M






Figure 2-71
The 765M/775M
routers.

Cisco 765M



Cisco 775M



transfer. It is based on the Intel 386 processor and comes with 1.5 MB, expandable to 2 MB over DRAM. The on-board NVRAM is 16 KB with 1 MB of flash memory. It can support up to 1,500 users and is available worldwide.

■ 766M/776M


Shown in Figure 2-72, these models include an internal NT1 device for connectivity. Additionally, these models have a second BRI port for external ISDN device connectivity or a second ISDN BRI line. It also includes two analog POTS RJ-11 ports for attaching phones, fax machines, and modems to share the ISDN BRI simultaneously with data. This model also supports provider supplemental services over ISDN such as call waiting, cancel call-waiting, call retrieve, call hold, three-way call conferencing, and call transfer. It is based on the Intel 386 processor and comes with 1.5 MB, expandable to 2 MB over DRAM. The on-board NVRAM is 16 KB with 1 MB of flash memory. It can support up to 1,500 users and is available in North America only.

The 700M family can act as the DHCP server for the LAN-attached devices, assigning the remote locations IP addresses to the attached workstations. The 700M family can also have its IP or IPX addresses assigned from the provider or central site network connection using Multilink Point-


Figure 2-72

The 766M/776M routers.

Cisco 766M



Cisco 776M



to-Point Protocol (MPPP). The ISDN BRI connection can dial on demand dynamically when it senses “interesting” traffic, as defined by the remote location network administrator. This feature is useful when one ISDN BRI B channel connects to one location and traffic is generated for a second location. The second B channel can be activated for the life of the interesting traffic and then be terminated. Also useful is setting FTP traffic as interesting to the router when transferring a large file to another location by bringing up the second B channel to increase bandwidth.



In typical configurations, many LAN workstation require access to another remote location. In many instances, the 700M is used as a connection to the Internet. Internet service providers (ISPs) typically provide only one Internet address for the location. The 700M uses a many-into-one feature called Port and Address Translation (PAT) to overcome this single address restriction. PAT is also used as a firewall function, protecting unknown resources from accessing the remote location and privileging internal devices to access the Internet. The access can include Web browsing, e-mail, or file transfer to devices on the remote LAN network.

As described, the 700M family is an access router. Its typical use is for occasional connectivity requirements from a remote location to another location. The location can be another remote office, the Internet, or a central office location.

CHAPTER

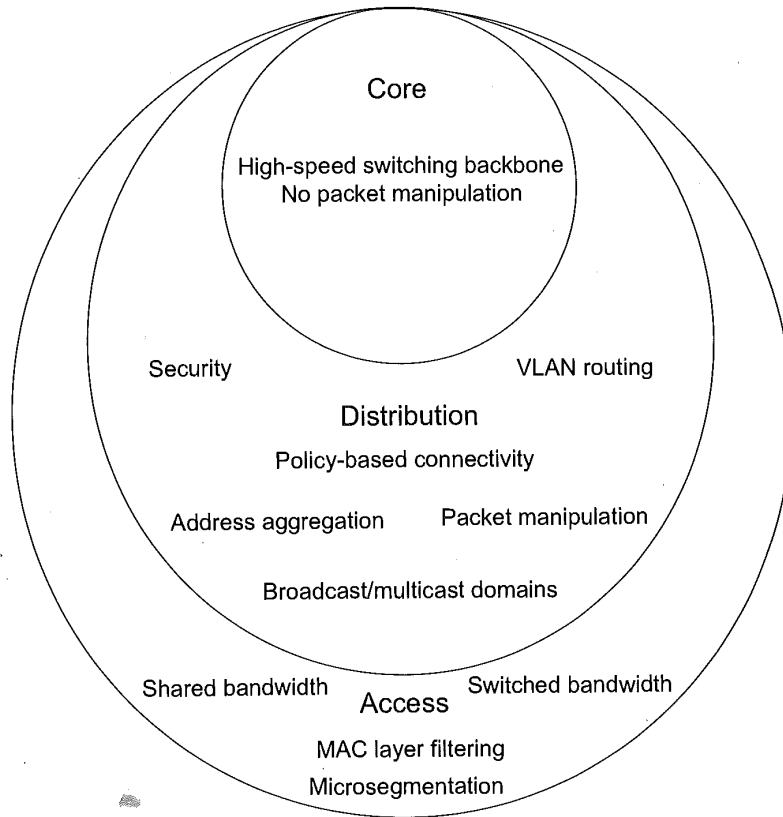
3

CISCO ROUTER NETWORK DESIGN



The hierarchical structure of the Cisco router network design model is based on the type of services provided at each layer. The notion of using layers creates a modular architecture, enabling growth and flexibility for new technologies at each layer. The Cisco hierarchical design model consists of three layers. Figure 3-1 diagrams the Cisco hierarchical design model.

Figure 3-1
The Cisco
Hierarchical Design
Model



The core layer provides the high-speed backbone for moving data between the other layers. This layer is geared towards the delivery of packets and not packet inspection or manipulation.

The distribution layer provides policy-based networking between the core and access layer. The distribution layer provides boundaries to the network topology and provides several services. These services are

- Address or area aggregation
- Departmental or workgroup access
- Broadcast/multicast domain definition
- *Virtual LAN* (VLAN) routing
- Any media transitions that need to occur
- Security

The access layer is the edge of the network. Being on the edge, the access layer is the entry point to the network for the end user community. Devices participating in the access layer can perform the following functions:

- Shared bandwidth
- Switched bandwidth
- MAC layer filtering
- Microsegmentation

It is important to remember that the Cisco hierarchical design model addresses the functional services of a network. The different layers described can be found in routers or switches, and each device can partake in the functions of more than one layer. Separation of functional layers is not mandatory, however. Maintaining a hierarchical design fosters a network optimized for performance and management.

The Network Infrastructure Lifecycle

Every corporation has a network infrastructure in place as the framework supporting the business processes. Just as applications and systems have lifecycles, so does a network infrastructure. This section highlights a network infrastructure lifecycle that can be used as a general guideline for designing and implementing Cisco-based networks.

Executive Corporate Vision

Corporate organizational restructuring through regional consolidation or through business group integration will certainly have an effect on the network infrastructure. Aligning the corporate vision with the business directives builds the foundation for the network infrastructure.

Gather Network Infrastructure Information

This involves research and discovery of the current network WAN topology as well as corporate and branch office LAN topologies. A full understanding of end-to-end network configuration is required. Additionally, bandwidth allocations and usage costs must be determined to provide the complete picture.

Determine Current Network Requirements

Communication protocols, client/server architectures, e-mail, distributed processing, Internet and intranet, voice and video, each has its own unique characteristics and can place demands on the network. These demands have to be recognized and understood for planning an enterprise-wide solution. The result from this study is a network profile for each business process and the network itself.

Assess Current Network Operational Processes

Network operational processes involve not just daily troubleshooting, but the other disciplines of network management: inventory, change, configuration, fault, security, capacity/performance, and accounting. Documenting the processes in place today will assist in evaluating the current baseline of service provided and identify areas that may need reengineering to meet the changing business requirements.

Research Plans for New Applications

The effect of new applications on network characteristics must be discovered prior to business groups moving into development, testing, and

production. Desktop video conferencing and voice communications along with data traffic require up-front knowledge to reengineer a network. Business group surveys and interviews, along with each group's strategic plan, will provide input to creating a requirements matrix.

Identify Networking Technologies

The selection of the appropriate technologies and how they can be of use in meeting current and future networking requirements relies on vendor offerings and their support structure. Paramount to this success is the partnership with and management of the vendors through an agreed-on working relationship.

Define a Flexible Strategic/Tactical Plan

The strategic plan in today's fast-paced changing technology environment requires flexibility. A successful strategic plan requires business continuity through tactical choices. The strategic plan must demonstrate networking needs in relation to business processes, both current and future.

Developing an Implementation Plan

This is the most visible of all the previous objectives. The planning and research performed prior to this can be for naught if the implementation does not protect current business processes from unscheduled outages. This must meet current business requirements and demands while migrating the network infrastructure to the strategic/tactical design. The perception to the business community must be business as usual.

Management and Review

The effectiveness of the new infrastructure is achieved through management and review. Reports highlighting the network health measured against expected service levels based on the strategic/tactical plan and design reflect the ability of the network to meet business objectives. The tools and analysis used here provide the basis for future network infrastructures.

Design Criteria

In planning for your network design, many criteria must be considered. These criteria are based on the current network design and performance requirements, as measured against the business direction compared to internetworking design trends. The trends of internetworking design affect the four distinct components of an enterprise internetwork:

Local area networks (LANs): These are networks within a single location that connect local end users to the services provided by the entire enterprise network.

Campus networks: These are networks within a small geographic area, interconnecting the buildings that make up the corporate or business entity for the area.

Wide area networks (WANs): These networks span large geographic areas and interconnect campus networks.

Remote networks: These types of networks connect branch offices, mobile users, or telecommuters to a campus or the Internet.

Figure 3-2 illustrates today's typical enterprise-wide corporate network topology.

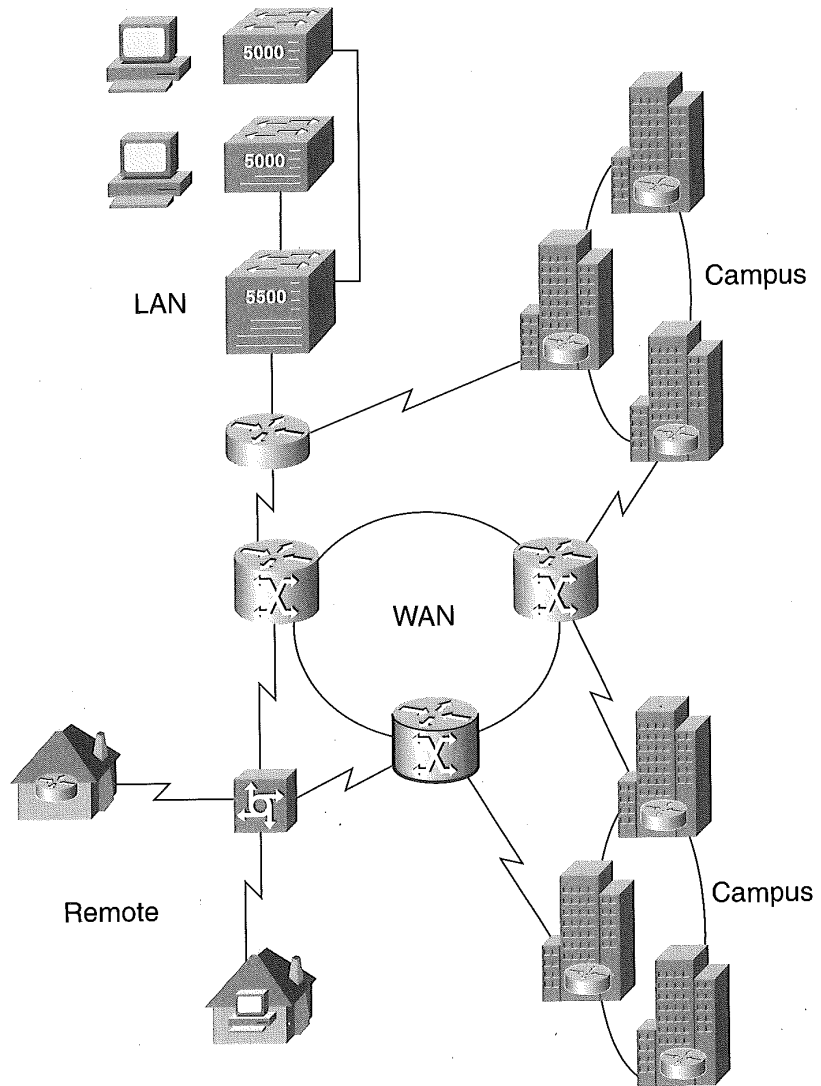
The Current LAN/Campus Network Trend

LANs and campus networks are grouped together for the simple reason that they share many of the same networking issues and requirements. Depending on the technologies used, a LAN can be focused within a building or span buildings. The spanning of a LAN makes up the campus network. Figure 3-3 diagrams a LAN/campus network topology.

Campus networks are a hybrid of LANs and WANs. From LAN/WAN technologies, campus networks use


- Ethernet
- Token Ring
- Fiber Distributed Data Interface (FDDI)
- Fast Ethernet
- Gigabit Ethernet
- Asynchronous Transfer Mode (ATM)
- T1/T3 networks
- Frame Relay

Figure 3-2
The enterprise-wide internetwork components



Two LAN technologies that serve to increase throughput and flexibility for LAN design are Layer 2 and Layer 3 switching. In short, Layer 2 switching occurs at the data link layer of the OSI Reference model and Layer 3 switching occurs at the Network layer of the OSI Reference model. Both switching algorithms increase performance by providing higher bandwidth

Figure 3-3
LAN/Campus
Network Topology



to attached workgroups, local servers, and workstations. The switches replace LAN hubs and concentrators in the wiring closets of the building.

The ability to switch end user traffic between ports on the device has enabled the concept of *Virtual LANs* (VLANs). Defining VLANs on the physical LAN enables logical groupings of end user segments or workstations. This enables traffic specific to this VLAN grouping to remain on this virtual LAN, rather than use bandwidth on LAN segments that are not interested in the grouped traffic. For example, the Finance VLAN traffic does not affect the Engineering VLAN traffic. Table 3-1 lists the important technologies affecting LAN and Campus network design.

WAN Design Trends

Routers are typically the connection points to WANs. Being at this juncture, the routers have become an important decision point for the delivery of traffic. With the advent of switching, the routers are slowly moving away from being the WAN device. The WAN services are now being handled by switches with three types of switching technologies:

- circuit switching
- packet switching
- cell switching

Table 3-1
Key LAN and
Campus Network
Technologies

| | |
|----------------------------|---|
| Routing technologies | Routing has long been the basis for creating internetworks. For use in a LAN/campus environment, routing can be combined with Layer 3 switching, which may also replace the entire function of a router. |
| LAN switching technologies | |
| Ethernet switching | Ethernet switching is Layer 2 switching, which can enable improved performance through dedicated Ethernet segments for each connection. |
| Token Ring switching | Token Ring switching is also Layer 2 switching. Switching Token Ring segments offers the same functionality as Ethernet switching. Token Ring switching operates as either a transparent bridge or a source-route bridge. |
| ATM switching technologies | ATM switching offers high-speed switching technology that integrates voice, video, and data. Its operation is similar to LAN-switching technologies for data operations. |

Circuit switching provides dedicated bandwidth, while packet switching enables the efficient use of bandwidth with flexibility to service multiple requirements. Cell switching combines the best of both circuit- and packet-switched networks. ATM is the leading cell-switched technology used in WANs today.

Because the WAN links end up servicing all traffic from one location to another, it is important that the bandwidth and performance be optimized. The optimization is due in part to the explosive growth of remote site connectivity, enhanced application architectures such as client/server and intranets, and the recent development of consolidating servers to a centralized location to ease administration and management. These factors have reversed the rules for traffic profiles from that of 80 percent LAN and 20 percent WAN to 80 percent WAN and 20 percent LAN.

This flip-flop of traffic characteristics has not only increased the requirement for WAN traffic optimization, but also for path redundancy, dial backup, and *quality of service* (QoS) to ensure application service levels over the WAN. The technologies available today that enable the effective and efficient use of WANs are summarized in Table 3-2. Coming on the horizon are such technologies as *Digital Subscriber Line* (DSL), *Low-Earth Orbit* (LEO) satellites, and advanced wireless technologies.

Remote Network Trends

Branch offices, telecommuters, and mobile users constitute remote networks. Some of these may use dial-up solutions with ISDN or analog modems. Others may require dedicated lines, allowing access to the WAN 24 hours a day, seven days a week (24x7). A study of the user's business requirements will dictate the type of connection for these remote locations.

Using ISDN and vendor functionality, a remote location can be serviced with 128 Kbps of bandwidth to the WAN only when traffic is destined out of the remote location. Analysis of the ISDN dial-up cost, based on uptime to the WAN as compared to the cost of a dedicated line to the WAN, must be determined for each location. This analysis will provide a break-even point on temporary versus dedicated WAN connectivity. Any of the various technologies discussed for the WAN may be well suited for remote network connectivity.

Table 3-2
WAN Technologies
and Use

| WAN Technology | Typical Uses |
|--|---|
| Analog modem | These are typically used for temporary dial-up connections or for backup of another type of link. The bandwidth is typically 9.6 Kbps to 56 Kbps. |
| Leased line | Leased lines have been the traditional technology for implementing WANs. These are links "leased" from communications services companies for exclusive use by your corporation. |
| Integrated Services Digital Network (ISDN) | An ISDN is a dial-up solution for temporary access to the WAN but adds the advantage of supporting voice/video/fax on the same physical connection. As a WAN technology, ISDN is typically used for dial-backup support at 56, 64, or 128 Kbps bandwidth. |
| Frame Relay | This is a distance-insensitive telco charge, thereby making it very cost-effective. It is used in both private and carrier-provided networks and most recently is being used to carry voice/video/fax/data. |
| Switched Multimegabit Data Service (SMDS) | SMDS provides high-speed, high-performance connections across public data networks. It can also be deployed in Metropolitan Area Networks (MANs). It is typically run at 45 Mbps bandwidth. |
| X.25 | X.25 can provide a reliable WAN circuit, but it does not provide the high bandwidth requirements as a backbone technology. |
| WAN ATM | This is used as the high-bandwidth backbone for supporting multiservice requirements. The ATM architecture supports multiple QoS classes for differing application requirements, delay and loss. |
| Packet over SONET (POS) | POS is an oncoming technology that transports IP packets encapsulated in SONET or SDH frames. POS meets the high bandwidth capabilities of ATM and through vendor implementations supports QoS. |

Application Availability Versus Cost-Effectiveness

It is the job of the network to connect end users with their applications. If the network is not available, then the end users are not working and the company loses money. Application availability is driven by the importance of the application to the business. This factor is then compared against the cost of providing application availability using the following:

- Redundant lines for alternate paths
- Dial-backup connectivity
- Redundant devices with redundant power supplies for connecting the end users
- Onsite or remote technical support
- Network management reach into the network for troubleshooting
- Disaster recovery connectivity of remote locations to the disaster recovery center

Designing an internetwork therefore has the main objective of providing availability and service balanced with acceptable costs for providing the service. The costs are generally dominated by three elements of supporting a network infrastructure:

- The number and location of hosts, servers, terminals, and other devices accessing the network; the traffic generated by these devices and the service levels required to meet the business needs.
- The reliability of the network infrastructure and traffic throughput that inherently affects availability and performance, thereby placing constraints on meeting the service levels required.
- The capability of the network equipment to interoperate, the topology of the network, the capacity of the LAN and WAN media, and the service required by the packets all affect the cost and availability factor.

The ultimate goal is to minimize the cost of these elements while at the same time delivering higher availability. The *total cost of ownership* (TCO), however, is dependent on understanding the application profiles.

Application Profile

Each application that drives a business network has a profile. Some profiles are based on corporate department requirements and others may be a directive for the entire company. A full understanding of the underlying architecture of the application and its use of the network is required for creating an application profile. Three basic components drive a network profile:

- Response time
- Throughput
- Reliability

Response time is a perceived result by the end user and a measured function of the network engineer. From a user standpoint, it is the reduced "think-time" of interactive applications that mandates acceptable response time. However, a network design that improves response time is relative to what the end user has perceived as normal response time.

A network engineer will break down the components that make up the response time into the following components:

- host time
- network time

The difference between the two is that host time is application processing, be this disk access to retrieve data or analysis of data. Network time is the transit time, as measured from leaving the host to the network interface of the end user device. Host time is then again computed on the workstation. Typically, host time on a workstation is based on presentation to the end user. Online interactive applications require low response times. These applications are usually referred to as time-sensitive applications.

Applications that rely on the delivery of large amounts of data are termed throughput-intensive applications. Typically, these applications perform file transfers. They require efficient throughput, yet many of these applications also depend on the delivery of the data within a time window. This is where they can adversely affect interactive application response times due to their throughput.

Reliability is often referred to as uptime. Applications requiring a high reliability inherently require high accessibility and availability. This in turn requires hardware and topology redundancy not only on the network side, but also on the application host or server side. The importance of the function served by the application is weighed against the cost of downtime incurred by the business. The higher the cost of downtime, the higher the requirement for reliability.

Creating an application becomes paramount in understanding the needs of a network design. Application profiles are assessed through exercising some or all of the following methods:

- *Profile the user community:* Determine corporate versus departmental internetworking requirements by separating common applications from specific applications for each community. If possible, develop the application flow from the end user to the host/server for each common and specific application. Using network management tools, gather network traffic profiles to parallel the user community.
- *Interviews, focus groups, and surveys:* Using these methods, insights into current perceptions and planned requirements are discovered.

This process is key for developing the current baseline of the network, in addition to coalescing information about planned requirements shared by independent departments. Data gathered here in combination with the community profiles is used for developing the new network design.

- *Design testing:* This is the proof-of-concept stage for the resulting design. Using simulated testing methods or real-time lab environments, the design is measured against the requirements for response-time, throughput, and reliability.

Cost-Efficiency

The network is now an asset to all corporations. As such, investment into the network must be viewed as a *total cost of ownership* (TCO). These costs are not only equipment investments but also include the following:

- *Total cost of equipment:* This includes not only hardware, but software, installation costs, maintenance costs, and upgrade costs.
- *Cost of performance:* This is a variable which measures improved network performance and reliability against the increase of business conducted. The ratio between the two determines the effectiveness of the investment.
- *Installation cost:* The physical cabling infrastructure to support the new design becomes a large one-time investment cost. A physical cabling infrastructure should be implemented that meets current and future networking technology requirements.
- *Growth costs:* These costs can be reduced by implementing technologies today that can follow the direction of technologies tomorrow.
- *Administrative and Support:* These limit the complexity of the internetwork design. The more complicated they are, the higher the costs for training, administration, management, and maintenance.
- *Cost of downtime:* Analyze the cost of limited, reduced, or inaccessible application hosts, servers, and databases. A high downtime cost may require a redundant design.
- *Opportunity costs:* Network design proposals should provide a minimum of two designs with a list of pros and cons to each design. Opportunity costs are the costs that may be realized by not choosing a

design option. These costs are measured more in a negative way; not moving to a new technology may result in competitive disadvantage, higher productivity costs, and poor performance.

- *Investment protection:* The current network infrastructure is often salvaged due to the large investment in cabling, network equipment, hosts, and servers. However, for most networks, investment costs are recovered within three years. Understand the cycle of cost recovery at your corporation. Apply this understanding to the design as a corporate advantage in the design proposal.

Keep in mind that the objective of any network design is the delicate balance of meeting business and application requirements while minimizing the cost to meet the objective.

Network Devices and Capabilities

The phenomenal growth of internetworks has predicated the move from bridges to routers and now switches. Four basic devices are used in building an internetwork. Understanding the functions of each is important in determining the network design. These four devices are

- Hubs
- Bridges
- Routers
- Switches

Hubs are often called concentrators and made centralized LAN topologies possible. All the LAN devices are connected to the hub. The hub essentially regenerates the signal received from one port to another acting as a repeater. These devices operate at the physical layer (Layer 1) of the OSI Reference Model.

Bridges connect autonomous LAN segments together as a single network and operate at the data link layer (Layer 2) of the OSI Reference Model. These devices use the *Media Access Control* (MAC) address of the end station for making a decision forwarding the packet. Bridges are protocol-independent.

Routers performing a routing function operate at the network layer (Layer 3) of the OSI Reference Model. These devices connect different networks and separate broadcast domains. Routers are protocol-dependent.

Switches were at first advanced multiport bridges with the capability to separate collision domains. Layer 2 switches that enhance performance and functionality through virtual LANs have replaced hubs. The second incarnation of switches enables them to perform Layer 3 routing decisions, thereby performing the function of a router.

Bridging and Routing

Bridging for this discussion is concerned with transparent bridging. This is opposed to *source-route bridging* (SRB), which is closer to routing than bridging. Bridging occurs at the MAC sublayer of the IEEE 802.3/802.5 standard applied to the data link layer of the OSI Reference Model. Routing takes place at the Network layer of the OSI Reference Model. Bridging views the network as a single logical network with one hop to reach the destination. Routing enables multiple hops to and between multiple networks. This leads to four distinct differences between the routing and bridging:

- The data-link packet header does not contain the same information fields as network layer packets.
- Bridges do not use handshaking protocols to establish connections, while network layer devices do utilize them.
- Bridges do not reorder packets from the same source, while network layer protocols expect reordering due to fragmentation.
- Bridges use MAC addresses for end node identification. Network layer devices, such as routers, use a network layer address associated with the wire to which the device is attached.

Although these differences exist between bridging and routing, bridging occasionally may be required or preferred over routing and vice versa.

The advantage of bridging over routing are as follows:

- Transparent bridges are self-learning and therefore require minimal, if any, configuration. Routing requires definitions for each interface for the assignment of a network address. These network addresses must be unique within the network.
- Bridging has less overhead for handling packets than does routing.
- Bridging is protocol-independent, while routing is protocol-dependent.
- Bridging will forward all LAN protocols. Routing only uses network layer information and therefore can only route packets.

In contrast, routing has the following advantage over bridging:

- Routing allows the best path to be chosen between source and destination. Bridging is limited to a specific path.
- Routing is a result of keeping updated, complete network topology information in routing tables on every routing node. Bridging maintains a table of devices connecting through the interface. This causes bridges to learn the network slower than routing, thereby enabling routing to provide a higher level of service.
- Routing uses network layer addressing, which enables a routing device to group the addresses into areas or domains, creating a hierarchical address structure. This leads to an unlimited amount of supported end nodes. Bridging devices maintain data link layer MAC addresses; therefore, they cannot be grouped and result in a limited number of supported end nodes.
- Routing devices will block broadcast storms from being propagated to all interfaces. Bridging spans the physical LAN segment to multiple segments and therefore forwards a broadcast to all attached LAN segments.
- Routing devices will fragment large packets to the smallest packet size for the selected route and then reassemble the packet to the original size for delivery to the end device. Bridges drop packets that are too large to send on the LAN segment without notification to the sending device.
- Routing devices notify transmitting end stations to slow down the transmission of data (congestion feedback) when the network itself becomes congested. Bridging devices do not possess that capability.

The general rule of thumb in deciding to route or bridge is to bridge only when needed. Route whenever possible.

Switching

The process of switching is the movement of packets from the receiving interface to a destination interface. Layer 2 switching uses the MAC address found within the frame. Layer 3 switching uses the network address found within the frame.

Layer 2 switching is essentially transparent bridging. A table is kept within the switching device for mapping the MAC address to the associated interface. The table is built by examining the source MAC address of each frame as it enters the interface. The switching function occurs when the


destination MAC address is examined and compared against the switching table. If a match is found, the frame is sent out of the corresponding interface. A frame that contains a destination MAC address not found in the switching table is broadcast out to all interfaces on the switching device. The returned frame will allow the switching device to learn the interface and therefore place the MAC address in the switching table.

MAC addresses are predetermined by the manufacturers of the *network interface cards* (NICs). These cards have unique manufacturer codes assigned by the IEEE with a unique identifier assigned by the manufacturer. This method virtually ensures unique MAC addresses. These manufacturer addresses are often referred to as *burned-in addresses* (BIA) or *universally administered addresses* (UAA). Some vendors, however, allow the UAA to be overridden with a *locally administered address* (LAA). Layer 2 switched networks are inherently considered flat networks.

In contrast, Layer 3 switching is essentially the function of a router. Layer 3 switching devices build a table similar to the Layer 2 switching table. Except in the case of the Layer 3 switching table, the entries are mapping network-layer addresses to interfaces. Since the network-layer addresses are based on assigning a logical connection to the physical network, a hierarchical topology is created with Layer 3 switching. As packets enter an interface on a Layer 3 switch, the source network-layer address is stored in a table that cross-references the network-layer address with the interface. Layer 3 switches carry with them the function of separating broadcast domains and network topology tables for determining optimal paths.

Combining Layer 2 and Layer 3 switching within a single device reduces the burden on a router to route the packet from one location to another, as shown in Figure 3-4. Switching therefore increases throughput due to the

Figure 3-4
Combined Layer 2
and Layer 3
switching and the
Cisco router



decisions being done in silicon, reduces central processing unit (CPU) overhead on the router, and eliminates hops between the source and destination device.

Backbone Considerations

The network backbone is the core of the three layer hierarchical model. Many factors affect the performance of the backbone:

- Path optimization
- Traffic prioritization
- Load balancing
- Alternate paths
- Switched access
- Encapsulation (Tunneling)

Path optimization is generally a function of a router that occurs using the routing table created by the network layer protocols. Cisco routers support all of the widely implemented IP routing protocols. These include

- *Open Shortest Path First (OSPF)*
- RIP
- IGRP
- EIGRP
- *Border Gateway Protocol (BGP)*
- *Exterior Gateway Protocol (EGP)*
- HELLO

Each of these routing protocols calculates the optimal path from the information provided within the routing tables. The calculation is based on metrics such as bandwidth, delay, load, and hops. When changes occur in the network, the routing tables are updated throughout all the routers within the network. The process of all the routers updating their tables and recalculating the optimal paths is called convergence. With each new generation of IP routing protocols, the convergence time is reduced. Currently, the IP routing calls with the smallest convergence times are Cisco proprietary routing protocols, IGRP and EIGRP.

Traffic prioritization enables the router to prioritize packets on the interface queue for delivery. This allows time-sensitive and mission-critical traf-

fic to take precedence over throughput-sensitive traffic. Cisco routers employ three types of traffic prioritization:

- Priority queuing
- Custom queuing
- Weighted-fair queuing

Priority queuing is the simplest form of traffic prioritization. It is designed primarily for low-speed links. The traffic under priority queuing is classified based on criteria that includes protocol and subprotocol types. The criteria profile is then assigned to a one-of-four output queuing. These queues are high, medium, normal, and low. In IP-based networks, the IP *type of service* (TOS) feature and Cisco IOS software capability to prioritize IBM logical unit traffic enable priority queuing for intraprotocol prioritization.

Custom queuing answers a fairness problem that arises with priority queuing. With priority queuing, low-priority queues may receive minimal service, if any. Custom queuing addresses this problem by reserving bandwidth for a particular type of traffic. Cisco custom queuing therefore allows the prioritization of multiprotocol traffic over a single link. For example, the greater the reserved bandwidth for a particular protocol, the more service is received. This provides a minimal level of service to all traffic over a shared media. The exception to this is underutilization of the reserved bandwidth. If traffic is not consuming the reserved bandwidth percentage, then the remaining percentage of reserved bandwidth will be shared by the other protocols. Custom queuing can use up to 16 queues, which are serviced sequentially until the configured byte count has been sent or the queue is empty.

Weighted-fair queuing uses an algorithm similar to time-division multiplexing. Each session over an interface is placed into a queue and allocated a slice of time for transmitting over the shared media. The process occurs in a round-robin fashion. Allowing each session to default to the same weighting parameters ensures that each session will receive a fair share of the bandwidth. This use of weighting protects time-sensitive traffic by ensuring available bandwidth and therefore consistent response times during heavy traffic loads.

The weighted-fair algorithm identifies the data streams over an interface dynamically. Because the algorithm is based on separating the data streams into logical queues, it cannot discern the requirements of different conversations that may occur over the session. This is an important point when considering queuing methods for protecting IBM SNA traffic.

Weighted-fair queuing becomes a disadvantage for SNA traffic when the SNA traffic is encapsulated in DLSw+ or RSRB.

The differences between the three queuing methods are dependent on the needs of the network. However, from an administrative point of view, weighted-fair queuing is far easier due to it being a dynamically-built queue versus priority and custom queuing, which both require the definitions of access lists, pre-allocated bandwidth, and predefined priorities.

Load balancing for IP traffic occurs with two to four paths to the destination network. It is not necessary for these paths to be of equal cost. The load balancing of IP traffic may occur on a per-packet basis or a per-destination basis. Bridged traffic over multiple serial links becomes balanced by employing a Cisco IOS software feature called circuit groups. This feature logically groups the multiple links as a single link.

Redundancy is a major design criterion for mission-critical processes. The use of alternate paths not only requires alternate links but requires terminating these links in different routers. Alternate paths are only valuable when a single point of failure is avoided.

Recovery of dedicated leased connections is mandatory for ensuring availability and service. This function is often termed switch access or switched connection; however, it does not relate to the Layer 2 or Layer 3 switching function. Switched access calls for the instantaneous recovery of WAN connectivity due to an outage on the dedicated leased line.

Switched access is also used to supplement bandwidth requirements using a Cisco IOS software feature called *bandwidth on demand* (BOD), which uses *dial-on-demand routing* (DDR). Using DDR along with the dedicated leased WAN connection, a remote location can send large mounts of traffic in a smaller time frame.

Encapsulation techniques are used for transporting non-routable protocols. IBM's SNA is a non-routable protocol. Encapsulation techniques are also used when the design calls for a single protocol backbone. These techniques are also referred to as tunneling.

Distributed Services

Within the router network, services can be distributed for maximizing bandwidth utilization, routing domains, and policy networking. The Cisco IOS software supports these distributed services through the following:

- Effective backbone bandwidth management
- Area and service filtering

- Policy-based distribution
- Gateway services
- Route redistribution
- Media translation

Preserving valuable backbone bandwidth is accomplished using the following features of Cisco IOS software:

- Adjusting priority output queue lengths, so overflows are minimized.
- Adjust routing metrics such as bandwidth and delay to facilitate control over path selection.
- Terminate local polling, acknowledgment, and discovery frames at the router using proxy services to minimize the high volume of small-packet traffic over the WAN.

Traffic filtering provides policy-based access control into the backbone from the distribution layer. The access control is based on area or service. Typically, we see the use of service access controls as a means for limiting an application service to a particular segment on the router. Traffic filtering is based on Cisco IOS software access control lists. These access control lists can affect inbound and outbound traffic of a specific interface or interfaces by being either permitted or denied.

Policy-based networking is a set of rules that determine the end-to-end distribution of traffic to the backbone. Policies can be defined to affect a specific department, protocol, or corporate policy for bandwidth management. The CiscoAssure initiative is a policy-based direction that enables the various network equipment to work together to ensure end-to-end policies.

Gateway functions of the router enable different versions of the same networking protocol to the internetwork. An example of this is connecting a DECnet Phase V network with a DECnet Phase IV network. These DECnet versions have implemented different addressing schemes. Cisco IOS within the router performs as an *address translation gateway* (ATG) for transporting the traffic between the two networks. Another example is AppleTalk's translational routing between different versions of AppleTalk. Route redistribution enables multiple IP routing protocols to interoperate through the redistribution of routing tables between the two IP routing protocols within the same router.

There are times in corporate networks when communication between different media is a requirement. This is seen more and more with the expansion of networks and newer technologies. For the most part, media translation occurs between Ethernet frames and Token Ring frames. The

translation is not a one-for-one since an Ethernet frame does not use many of the fields used in a Token Ring frame. An additional translation that can be observed is one from IBM SDLC to Logical Link Control 2 (LLC2) frames. This enables serial-attached IBM SDLC connections to access LAN attached devices.

Local Services

At the local access layer of the three-layer model, features provided by the Cisco IOS within the router provide added management and control over access to the distribution layer. These features are

- Value-added network addressing
- Network segmentation
- Broadcast and multicast capabilities
- Naming, proxy, and local cache capabilities
- Media access security
- Router discovery

The discovery of servers and other services may sometimes cause broadcasts within the LAN. A feature on Cisco IOS software directs these requests to specific network-layer addresses. This feature is called helper addressing. Using this feature limits the broadcast to only segments of the helper addresses defined for that service. This is best used when protocols such as Novell IPX or DHCP typically search the entire network for a server using broadcast messages. Helper addresses thereby preserve bandwidth on segments that do not connect the server requested.

Network congestion is typically a result of a poorly designed network. Congestion is manageable by segmenting networks into smaller, more manageable pieces. Using multiple IP subnets, DECnet areas and AppleTalk zones further segment the network so that traffic belonging to the segment remains on the segment. Virtual LANs further enhance this concept by spanning the segmentation between network equipment.

Although routers control data link (MAC address) broadcasts, they allow network layer (Layer 3) broadcasts. Layer 3 broadcasts are often used for locating servers and services required by the host. The advent of video broadcasts has proliferated the use of multicast packets over a network. Cisco IOS does its best in reducing broadcast packets over IP networks through directed broadcasts to specific networks, rather than the entire network.

In addition, the Cisco IOS employs a spanning-tree technique when flooded broadcasts are recognized, minimizing excessive traffic but enabling the delivery of the broadcast to all networks. IP multicast traffic moves from a single source to multiple destinations. IP multicast is supported by a router running Cisco IOS with the *Internet Group Management Protocol* (IGMP) implemented. Using IGMP, the router can serve as a multicast distribution point, delivering packets to only segments that are members of the multicast group, ensuring loop-free paths, and eliminating duplicate multicast packets.

The Cisco IOS software contains many features for further reducing bandwidth utilization using naming, proxy, and local cache functions. The function drastically reduces discovery, polling, and searching characteristics of many of the popular protocols from the backbone. The following is a list of the features available with Cisco IOS that limits these types of traffic from the backbone:

- *Name services:* NetBIOS, DNS, and AppleTalk Name Binding Protocol
- *Proxy services:* NetBIOS, SNA XID/Test, polling, IP ARP, Novell ARP, AppleTalk NBP
- *Local Caching:* SRB RIF, IP ARP, DECnet, Novell IPX

Selecting Routing Protocol

Routing protocols are the transport of IP-based networks. The following are examples of routing protocols:

- *Routing Information Protocol* (RIP)
- *Routing Information Protocol 2* (RIP2)
- *Interior Gateway Routing Protocol* (IGRP)
- *Enhanced Interior Gateway Routing Protocol* (EIGRP)
- *Open Shortest Path First* (OSPF)
- *Intermediate System-Intermediate System* (IS-IS)

In selecting a routing protocol for the network, the characteristics of the application protocols and services must be taken into consideration. Network designs enabling a single routing protocol are best for network per-

formance, maintenance, and troubleshooting. Six characteristics of a network must be considered when selecting a routing protocol:

- Network topology
- Addressing and route summarization
- Route selection
- Convergence
- Network scalability
- Security

Network Topology

Routing protocols view the network topology in two ways: flat or hierarchical. The physical network topology consists of the connections of all the routers within the network. Flat routing topologies use network addressing to segregate the physical network into smaller interconnected flat networks. Examples of routing protocols that use a non-hierarchical flat logical topology are RIP, RIP2, IGRP, and EIGRP.

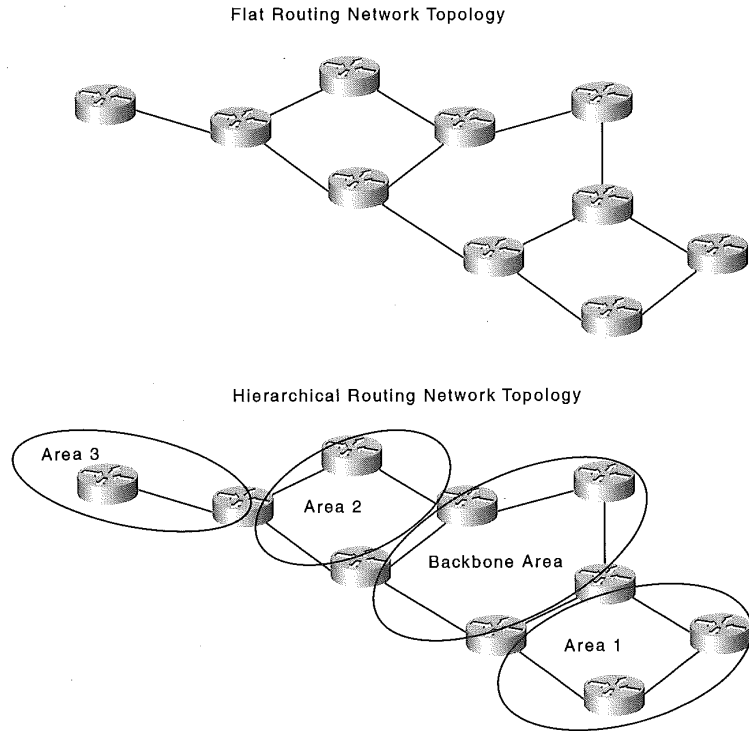
OSPF and IS-IS routing networks are hierarchical in design. As shown in Figure 3-5, hierarchical routing networks assign routers to a routing area or domain. The common area is considered the top of the hierarchy, which the other routing areas communicate through. Hierarchy routing topologies assign routers to areas. These areas are the routing network addresses, used for delivering data from one subnet to another. The areas are a logical grouping of contiguous networks and hosts. Each router maintains a topology map of its own area, but not of the whole network.

Addressing and Route Summarization

Some of the IP routing protocols have the capability to automatically summarize the routing information. Using summarization, the route table updates the flow between routers, which is greatly reduced, thereby saving bandwidth, router memory, and router CPU utilization. As shown in Figure 3-6, a network of 1,000 subnets must have 1,000 routes. Each of the routers within the network must therefore maintain a 1,000-route table. If we assume that the network is using a Class B addressing scheme with a subnet mask of 255.255.255.0, summarization reduces the number of routes within each router to 253. Three routes in each of the routers describe the

Figure 3-5

Flat versus hierarchical routing topologies.



path to the other subnets on the other routers, and 250 routes describe the subnets connected to each router.


Route Selection

In networks where high availability and redundancy are a requirement, the route selection algorithm of the routing protocol becomes an important factor in maintaining acceptable availability. Each of the routing protocols uses some type of metric to determine the best path between the source and the destination of a packet. The available metrics are combined to produce a “weight” or “cost” on the efficiency of the route.

Depending on the routing protocol in use, multiple paths of equal cost

Figure 3-6

The effect of route summarization.

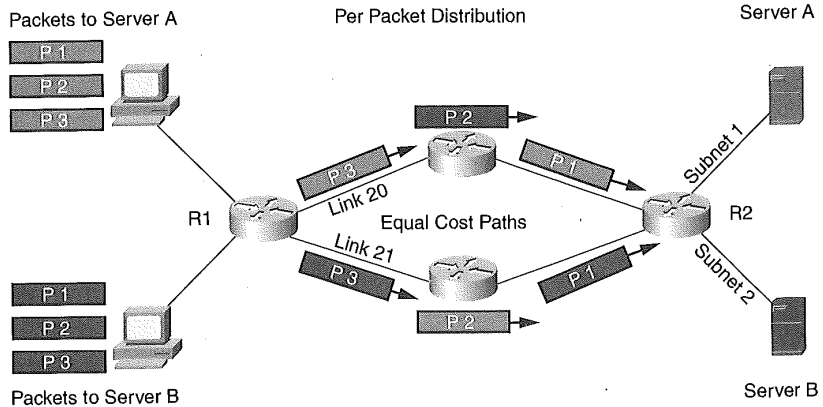


may provide load balancing between the source and destination, thereby spreading the load across the network. Some protocols like EIGRP can use unequal cost paths to load balance. This capability to load balance further improves the management of network bandwidth.

Load balancing over multiple paths is performed on a per-packet or per-destination basis. Per-packet distributes the load across the possible paths in proportion to the routing metrics of the paths. For equal cost paths, this results in a round-robin distribution. There is, however, the potential in a per-packet load balancing technique for the packets to be received out of order. Per-destination load balancing distributes the packets based on the destination over the multiple paths to the destination. For instance, as shown in Figure 3-7, packets destined for subnets attached to router R2 from router R1 use a round-robin technique based on the destination. Packets destined for subnet 1 flow over link 20, while packets destined for subnet 2 flow over link 21 versus the per-packet basis of alternating the packets for subnet 1 and subnet 2 over the two links.

Figure 3-7

Packet distribution using a round-robin technique



The Concept of Convergence

Convergence is the time it takes a router to recognize a network topology change, calculate the change within its own table, and then distribute the table to adjacent routers. The adjacent routers then perform the same functions. The total time it takes for the routers to begin using the new calculated route is called the convergence time. The time for convergence is critical for time-sensitive traffic. If a router takes too long to detect, recalculate, and then distribute the new route, the time-sensitive traffic may experience poor performance or the end nodes of the connection may drop.

In general, the concern with convergence is not with the addition of new links or subnets in the network. The concern is the failure of connectivity to the network. Routers recognize physical connection losses rapidly. The issue of long convergence time is the failure to detect poor connections within a reasonable amount of time. Poor connections such as line errors, high collision rates, and others require some customization of the router for detecting these types of problems faster.

Network Scalability

The capability of routing protocols to scale to a growing network is not so much a weakness of the protocol, but the critical resources of the router hardware. Routers require memory, CPU, and adequate bandwidth to properly service the network.

Routing tables and network topology are stored in router memory. Using a route summarization technique, as described earlier, reduces the memory requirement. In addition, routing protocols that use areas or domains in a hierarchical topology require the network design to use small areas, rather than large ones, to help in reducing the memory consumption.

Calculation of the routes is a CPU-intensive process. Through route summarization and the use of link-state routing protocols, the CPU utilization is greatly reduced since the number of routes needing re-computing is reduced.

Bandwidth on the connections to each router becomes a factor in not only scaling the network, but in convergence time. Routing protocols find information on neighbor routers for the purpose of receiving and sending routing table updates. The type of routing protocol in use will determine its affect on the bandwidth.

Distance-vector routing protocols such as RIP and IGRP send their routing tables at regular intervals. The distance-vector routing protocol waits for the time interval before sending its update, even when a network change has occurred. In stable networks, this type of updating mechanism wastes bandwidth yet protects the bandwidth from an excessive routing update load when a change has occurred. However, due to the periodic update mechanism, distance vector protocols tend to have a slow convergence time.

Link-state IP routing protocols such as OSPF and IS-IS address bandwidth wastefulness of distance-vector routing protocols and slow convergence times. Due to the complexity of providing this enhancement, however, link-state protocols are CPU-intensive and require higher memory utilization and bandwidth during convergence.

During network stability, link-state protocols take up minimal network bandwidth. After start-up and initial convergence, updates are sent to neighbors only when the network topology changes. During a recognized topology change, the router will flood its neighbors with the updates. This may cause excessive load on the bandwidth, CPU, and memory of each router. However, convergence time is lower than that of distance-vector protocol.

Cisco's proprietary routing protocol, EIGRP, is an advanced version of distance-vector protocols with properties of link-state protocols. EIGRP has taken many of the metrics for route calculation from distance-vector protocols. The advantages of link-state protocols are used for sending routing updates only when changes occur. Although EIGRP preserves CPU, memory, and bandwidth during a stable network environment, it does have high CPU, memory, and bandwidth requirements during convergence.

Table 3-3

Recommended
Number of Neigh-
bors per Router
with IP Routing
Protocols

| Routing Protocol | Neighbors per Router |
|----------------------------------|----------------------|
| Distance vector (RIP, IGRP) | 50 |
| Link state (OSPF, IS-IS) | 30 |
| Advanced distance vector (EIGRP) | 30 |

The convergence capability of the routing protocols and their effect on CPU, memory, and bandwidth has resulted in guidelines from Cisco on the number of neighbors that can effectively be supported. Table 3-3 lists the suggested neighbors for each protocol.

Security

Routing protocols can be used to provide a minimal level of security. Some of the security functions available on routing protocols are



- Filtering route advertisements
- Authentication

Using filtering, routing protocols can prohibit the advertisements of routes to neighbors, thereby protecting certain parts of the network. Some of the routing protocols authenticate their neighbor prior to engaging in routing table updates. Although this is protocol-specific and generally a weak form of security, it does protect unwanted connectivity from other networks using the same routing protocol.

CHAPTER

4

IP Routing Protocol Design



Routing is the process of moving packets from one network to another. A routing decision takes place at the source network device, which is a router. The decision is made based on metrics used for a particular routing protocol. Routing protocols can use some or all of the following metrics in determining the best route to a destination network:

- Path length
- Reliability
- Delay
- Bandwidth
- Load
- Communication cost

Path length is a measure of either a cost or a hop count. In link-state routing protocols, the cost is the sum of the costs associated with each link in the path. Distance-vector routing protocols assign a hop count to the path length, which measures the number of routers a packet traverses between the source and destination.

Reliability is typically the bit-error rate of a link connecting this router to a source or destination resource. For most of the routing protocols, the reliability of a link is assigned by the network engineer. Since it is arbitrary, it can be used to influence and create paths that are favorable over other paths.

The delay metric is an overall measurement of the time it takes for a packet to move through all the internetworked devices, links, and queues of each router. In addition, network congestion and the overall distance traveled between the source and destination are taken into consideration in evaluating the delay metric value. Because the delay value takes into account many different variables, it is an influential metric on the optimal path calculation.

Using bandwidth as a metric in optimal path calculations can be misleading. Although a bandwidth of 1.544 Mbps is greater than 56 Kbps, it may not be optimal due to the current utilization of the link or the load on the device on the receiving end of the link.

The load is a metric that assigns a value to a network resource based on the resource's overall utilization. This value is a composite of CPU utilization, the packets processed per second, and the disassembly or reassembly of packets, among other things. The monitoring of the device resources itself is an intensive process.

In some cases, communication lines are charged based on usage versus a flat monthly fee for public networks. For example, ISDN lines are charged based on usage time and the amount of data transmitted during that time.

In these instances, communication cost becomes an important factor in determining the optimal route.

In designing a routing protocol-based network, the routing algorithm should have the following characteristics built into the design:

- *Optimality:* This concerns using some or all of the metrics available for a routing protocol in order to calculate the optimal route. Different routing protocols may apply one metric as having a higher weight to the optimal route calculation than another. An understanding of this behavior is important in choosing the routing protocol.
- *Simplicity:* Although routing protocols themselves may be complicated, their implementation and operational support must be simplistic. Router overhead and the efficient use of router resources is important in maintaining a stable and reliable network.
- *Robustness:* Choose a routing algorithm that meets the requirements of the network design. In some cases, such as with small networks, a simplistic distance-vector routing protocol is sufficient. Large networks that require a hierarchical design require the capability of the routing protocol to scale to the size of the network without itself becoming a hindrance on the network.
- *Rapid Convergence:* The convergence time to recalculate and then use a new optimal path between a source and destination resource is paramount in meeting availability and service level requirements of a network.
- *Flexibility:* The algorithms employed by the selected routing protocol must be flexible and adapt to the changing dynamics of network resources and the network as a whole.

RIP, RIP2, and IGRP Network Design

RIP, RIP2, and IGRP are distance-vector-based routing protocols. Such routing protocols base the optimal route on the number of hops (devices, in other words) a packet must pass through to reach a destination.

Routing Information Protocol (RIP) was the first routing protocol algorithm for distributing, calculating, and managing available routes within a network. *Interior Gateway Routing Protocol* (IGRP) is a Cisco proprietary routing protocol algorithm using enhanced optimal route calculation. IGRP calculates optimal routes based on bandwidth, delay, reliability, and load. RIP2 is the second generation of RIP. RIP2 supports the Internet Protocol

Version 6 specification for 128-bit addressing, *variable-length subnet masks* (VLSM), and route summarization.


RIP, RIP2, and IGRP Topology

Distance-vector routing protocols use a flat network topology, as shown in Figure 4-1. Since these protocols are distance-vector-based routing algorithms, it is beneficial to minimize the number of hops between two destinations. This requires careful planning of the core, distribution, and access topology layers in planning the hierarchical service model. For most cases, when deploying distance-vector-based routing protocols, the service functions of the core, distribution, and access layers typically commingle within a single router.

RIP, RIP2, and IGRP Addressing and Summarization

In RIP and IGRP networks, the IP 16-bit addressing scheme of IP version 4 is supported. RIP2 supports both the IP version 4 16-bit and IP version 6

Figure 4-1
Flat topology of a distance-vector-based router network




128-bit addressing scheme. Additionally, RIP and IGRP support fixed subnet masks for a network. Every subnet address used in the RIP or IGRP network must use the same subnet masking. RIP2 using VLSM and the 128-bit addressing scheme allows for varied subnet masks of the router interface. This is because the RIP2 routing packet includes the subnet mask of the source and destination IP address. Summarization reduces the memory requirements on the router by keeping the routing table to a minimum.

RIP, RIP2, and IGRP Route Selection and Convergence

Both RIP and RIP2 base the optimal route selection on the number of hops. IGRP enhances this by incorporating bandwidth, delay, reliability, and load. Figure 4-2 illustrates the route selection difference between RIP, RIP2, and IGRP. RIP and IGRP use the first route within their routing tables as the optimal route for a destination network or subnet. Because IGRP uses bandwidth as a metric, the IGRP optimal route in the figure has more hops than the RIP optimal route. The Cisco IOS implementation of RIP provides for up to six entries on which the router can load-balance to a destination. IGRP will load-balance packets over equal-cost paths to a destination network or subnet. This load balancing occurs in a round-robin fashion.

Figure 4-2
Route selection differences between RIP, RIP2, and IGRP



Both RIP and IGRP build their tables and then transmit the entire routing table to adjacent routers. Each router in turn recalculates its table based on the information received from the sending router. Once this is completed, the router forwards its new table to adjacent routers.

Both RIP and IGRP also periodically send their routing tables to adjacent routers. RIP defaults to a 30-second interval for sending the routing table to adjacent routers. IGRP defaults to a 90-second interval for sending the routing table to adjacent routers. Both RIP and IGRP recalculate routing entries once recognizing a link outage or timeout to an adjacent router. The recalculated routing table is not forwarded to adjacent routers, however, until the update interval has been reached. The periodic updating of neighbor routers for topology changes causes excessive convergence time for the network to learn new optimal routes.

RIP2, however, addresses the periodic update problem by sending only the updated route entry at the time of the recalculation. Although this sounds much like a link-state protocol update, RIP2 still sends the entire table on a periodic basis. The capability of RIP2 to send an update at the time it is recalculated reduces the convergence time. RIP2 sends the entire routing table on a periodic basis, just as RIP and IGRP. However, the table is smaller due to the use of VLSM and route summarization. RIP2 will load-balance packets to a destination network or subnet over equal-cost paths.

RIP, RIP2, and IGRP Network Scalability

Time for the convergence of RIP, IGRP, and RIP2 networks is the single inhibitor to scaling these protocols to large networks. Convergence is not just a time factor, but also a CPU and memory issue on each router. These protocols recalculate the entire table during convergence, as opposed to just the affected route. Therefore, convergence becomes a CPU-intensive process, thereby reducing the capability of a router to provide service levels during convergence. Since these protocols send the entire table in a periodic time frame, they consume bandwidth, causing bandwidth constraints on an ongoing basis.

EIGRP Network Design

Enhanced Interior Gateway Protocol (EIGRP) is a proprietary routing protocol of Cisco Systems. EIGRP merges the best of distance-vector protocol characteristics with advantages of link-state protocol characteristics. In

addition, EIGRP uses *Diffusing Update Algorithm* (DUAL) for fast convergence and the further reduction of possible routing loops within the network. An advantage to using EIGRP over other routing protocols is its capability to support not only IP, but also Novell NetWare IPX and AppleTalk, thus simplifying network design and troubleshooting.

EIGRP Topology

EIGRP uses a non-hierarchical flat networking topology. It automatically summarizes subnet routers for networks directly connected to the router using the network number as the boundary. It has been found that the automatic summarization is sufficient for most IP networks.

EIGRP Addressing and Summarization

EIGRP supports *variable-length subnet masking* (VLSM). Defining an address space for use by an EIGRP is a primary step in developing the routing architecture. EIGRP support for VLSM is made possible by including the subnet mask assigned to the router interface in the EIGRP routing messages. VLSM is essentially the subnetting of a subnet (or a sub-subnet).

Using an appropriate addressing scheme, the size of the routing tables and convergence time can drastically be reduced through route summarization. EIGRP automatically summarizes the routes at network number boundaries. Figure 4-3 diagrams the use of route summarization.

The network engineer can configure route summarization at the interface level, however, using any bit-boundary of the address to further summarize the routing entries. The metric used in route summarization is the best route found for the routes used to determine the summarized route. The summary points to the NULL interface of the summarizing router. This makes the metric the cost of reaching the summarizing router.

EIGRP Route Selection

EIGRP uses the same metrics as IGRP. These values are bandwidth, delay, reliability, and load. The metric placed on a route using EIGRP defaults to the minimum bandwidth of each hop, plus a media-specific delay for each hop. The value for the metrics used in EIGRP are determined as follows:

- **Bandwidth:** EIGRP uses the default value for each interface to the value specified by the bandwidth interface command.

- **Delay:** The inherent delay associated with an interface. The delay metric can also be defined on an interface using the delay interface command.
- **Reliability:** A dynamically computed value averaged over five seconds. The reliability metric changes with each new weighted average.
- **Load:** A dynamically computed weighted average over five seconds. The load metric changes with each new weighted average.

EIGRP Convergence

EIGRP employs a *diffusing update algorithm* (DUAL) for calculating route computations. DUAL uses distance vector algorithms to determine loop-free efficient paths, selecting the best path for insertion into the routing table. DUAL, however, also determines the second-best optimal route for each entry; this route is termed a feasible successor. The feasible successor entry is used when the primary route becomes unavailable. Figure 4-4 illustrates the use of the feasible successor. Using this methodology of successor routes avoids a recalculation and therefore minimizes convergence time. Along with primary routes, EIGRP distributes the feasible successor entries to the neighboring routers.

Figure 4-3
EIGRP route
summarization
using VLSM

Subnet mask
255.255.255.240

10.22.33.16
10.22.33.32
10.22.33.48
10.22.33.64
10.22.33.80
10.22.33.96
10.22.33.112
10.22.33.128
10.22.33.144
10.22.33.160
10.22.33.176
10.22.33.192
10.22.33.208
10.22.33.224





Figure 4-4
EIGRP use of
a feasible successor


Stable Network

() Link cost
→ Cost to Subnet 8




Successor Discovery

() Link cost
→ Cost to Subnet 8



Stable Network

() Link cost
→ Cost to Subnet 8



EIGRP Scalability

Scalability is a function of memory, CPU, and bandwidth efficiencies. EIGRP is designed for optimizing these resources. Through route summarization, the routes advertised by neighbors are stored with minimal memory required. This enables an EIGRP network to expand without routing issues. Since EIGRP uses DUAL only, routes that are affected by a change are recomputed, and since EIGRP is based on the same metrics as IGRP, the computation CPU requirements are minimal. Because EIGRP only sends updates due to topology changes, bandwidth is preserved. Steady-state bandwidth utilization of EIGRP is minimal due to the use of EIGRP's HELLO protocol for maintaining adjacencies between neighbors.

EIGRP Security

Since EIGRP is a Cisco IOS proprietary routing protocol, it is available only on Cisco routers. Additionally, route filters and authentication can be specified to further limit accidental or malicious routing disruptions from unknown routers connecting to the network.

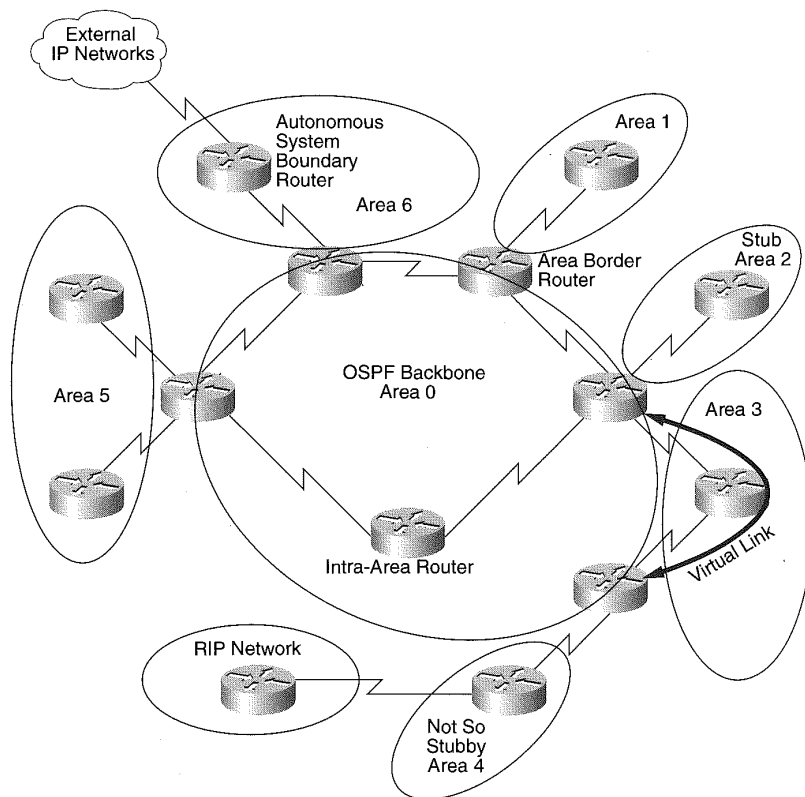
OSPF Network Design

Open Shortest Path First (OSPF) is a standards-based link-state routing protocol defined by the *Internet Engineering Task Force* (IETF) OSPF workgroup and published in *Request for Comment* (RFC) 1247. OSPF is based on an *autonomous system* (AS), which is defined by OSPF as a group of routers exchanging routing information using link-state protocol. OSPF is based on using a hierarchical networking topology. Defining the hierarchy requires planning to define boundaries that denote an OSPF area and address assignment.

OSPF Topology

OSPF defines its hierarchy based on areas. Figure 4-5 illustrates the OSPF hierarchy and various areas used to build and connect the OSPF network. An area is a common grouping of routers and their interfaces. OSPF has one single common area through which all other areas communicate. Due to the use of the OSPF algorithm and its demand on router resources, it is

Figure 4-5
OSPF network
hierarchy and areas



necessary to keep the number of routers at 50 or below per OSPF area. Areas with unreliable links will therefore require many recalculations and are best suited to operate within small areas.

The OSPF algorithm using a flooding technique for notifying neighbors of topology must change. The greater number of neighbors, the more CPU-intensive the topology becomes since the new route must be recalculated and forwarded to all attached neighbors. Cisco studies have resulted in a recommendation of no more than 60 neighbors per OSPF router.

The OSPF link-state algorithm calculates a change for each specified area defined on the router. Area routers are usually also *area border routers* (ABR). That is, they maintain and support OSPF routing tables for two OSPF areas. In general, there is a minimum of two areas for an ABR: the backbone area and one non-backbone area. The recommendation for OSPF

is to limit the number of supported areas in a router to three. This minimizes resources utilization for the calculation and distribution of link-state updates.

OSPF uses a designated router as the keeper of all the OSPF routes within a LAN. This reduces routing updates over a LAN, thereby preserving LAN media bandwidth. OSPF routers attached to the same LAN as the designated router request a route only if their own table does not have an entry for the destination resource. A backup designated router is also used for availability and redundancy. The recommendation is to have a designated and backup designated router supporting only one LAN. In addition, the designated and backup designated router should be the least CPU-intensive router on the LAN.

The OSPF backbone must be designed for stability and redundancy. A link failure that partitions the backbone will result in application outages, which leads to poor availability. The size of the backbone should be no more than 50 routers.

Routers within the OSPF backbone must be contiguous. This follows the concept of the hierarchy and maintains the traffic for backbone updates within the backbone area routers. However, OSPF offers the use of a virtual link for connecting two non-contiguous routers through a non-native area router. Using a virtual link, a partitioned backbone can be circumvented until the link failure causing the outage is corrected. Finally, reserve the media used for the OSPF backbone for routers to avoid instability and unrelated routing protocol traffic.

As with backbone areas, each OSPF area must be contiguous, and not only contiguous in design, but contiguous in the network address space. Using a contiguous address space makes route summarization possible. The routers of an area connecting the area to the OSPF backbone area are termed *area border routers* (ABR). For availability, it is deemed appropriate to have more than one ABR connecting the area to the backbone area.

Designing large-scale OSPF networks requires a review of the physical connectivity map between routers and the density of resources. Designing the network into geographic areas can be beneficial for simplifying implementation and operations but may not be beneficial for availability or performance. In general, smaller OSPF areas generate better performance and higher levels of availability than large OSPF areas.

OSPF Addressing and Summarization

Maximizing the address space in OSPF networks assists in reducing resource utilization and maximizes route summarization. A hierarchical

addressing scheme is the most effective means of designing an OSPF network. OSPF supports VLSM that lends itself to a hierarchical network address space specification. Using VLSM, route summarization is maximized at the backbone and ABR routers. Guidelines for defining an OSPF network for optimized route summarization are as follows:

- Define the network address scheme in subnet ranges for use in each contiguous area.
- Use VLSM addressing to maximize address space.
- Define the network address space for future growth to allow the splitting of an area.
- Design the network with the intention of adding new OSPF routers in the future.

Route summarization increases the stability of an OSPF network and keeps route changes within an area. Route summarization must be explicitly specified when working with OSPF networks on Cisco routers. The specification of router summarization requires the following information:

- Determine route information needed by the backbone about each area.
- Determine route information needed by an area for the backbone and other areas.

OSPF route summarization occurs in area border routers. Using VLSM, bit-boundary summarization is possible on network or subnet addresses within the area. Since OSPF route summarization is explicit, the network design must incorporate summarization definitions for each OSPF area border router.

OSPF areas offer four types of routing information:

Default: A default route for packets whose destination IP network or subnet cannot be found in the routing tables.

Intra-area routes: These are routes for network or subnets within a given area.

Interarea routes: This information provides areas with explicit network or subnet routers for networks or subnets within the OSPF autonomous system, but not within the area.

External routes: These are routes learned from the exchange of routing information between autonomous systems. This results in routes that are external to the OSPF autonomous system.

OSPF route information provides information on three types of OSPF areas: non-stub areas, stub areas, and stub areas without summaries. Stub

areas are OSPF areas that connect only to one other area and therefore are considered a stub off the hierarchy. A non-stub area is an OSPF area that provides connectivity to more than one OSPF area.

Non-stub area characteristics include the following:

- They store default routes, static routes, intra-area routes, interarea routes, and external routes.
- They have OSPF inter-area connectivity.
- Autonomous system border routers are used.
- Virtual links require non-stub areas.
- They are the most resource-intensive type of area.

Stub area characteristics are as follows:

- They build default, intra-area, and inter-area routes.
- They are most useful in areas containing one ABR.
- They may contain multiple area border routers to the same area.
- Virtual links cannot connect through stub areas.
- They cannot use autonomous system border routers.

Stub areas without summaries have the following characteristics:

- They serve as default and intra-area routers.
- They are recommended for single router connections to the backbone.

Table 4-1 lists the OSPF area types along with the routing information supported.

Table 4-1

OSPF Area Type
and Support
for Routing
Information
Routing Information
Type

| Area Type | Default | Intra-area | Interarea | External |
|------------------------|---------|------------|-----------|----------|
| Nonstub | Yes | Yes | Yes | Yes |
| Stub | Yes | Yes | Yes | No |
| Stub without summaries | Yes | Yes | No | No |

OSPF Route Selection

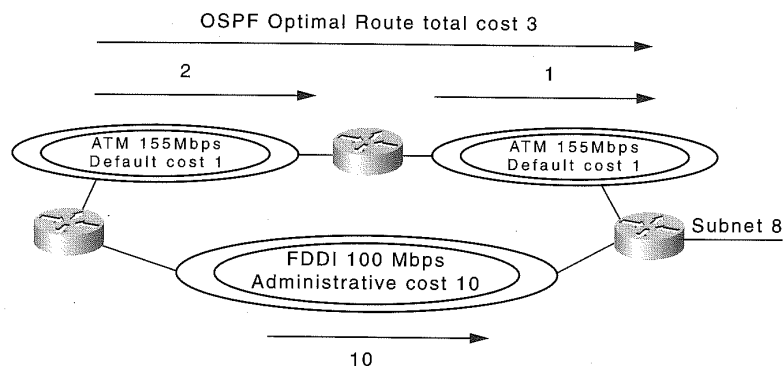
OSPF defaults route selection to the bandwidth metric. Under OSPF, the bandwidth metric is determined by the type of media being used and the cost is measured on the outgoing interface only. The bandwidth metric for a link is the inverse of the bandwidth supported by the media used for the link. The bandwidth metric is calibrated based on a metric of 1 for FDDI media. Figure 4-6 depicts an OSPF network and the applied bandwidth metric.

The total metric for a given route is the sum of all the bandwidth metric values of all the links used for the route. Media that support bandwidth greater than FDDI 100 Mbps default to the FDDI metric value of one. In a configuration where media types connecting the router are faster than FDDI, a manual cost greater than one must be applied to the FDDI link in order to favor the higher-speed media type. OSPF route summarization uses the metric of the best route found within the summarized routes as a metric value for the summarized entry.

OSPF external routes are defined as being either a type 1 or type 2 route. The metric for a type 1 external route is the sum of the internal OSPF metric and the external route metric. Type 2 external routes use only the metric of the external route. Type 1 external route metrics are more favorable in providing a truer metric for connecting to the external resource.

For single ABR OSPF areas, all traffic leaving the area flows through the single ABR. This is done by having the ABR exchange a default route with the other routers of the area. In multiple ABR OSPF areas, the traffic can leave either through the ABR closest to the source of the traffic or the ABR closest to the destination of the traffic. In this case, the ABRs exchange summarized routes with the other routers of the area.

Figure 4-6
OSPF route selection
using bandwidth
metrics



High-availability network design requires redundant paths and routers. Redundancy is useful when employing equal-cost paths to take advantage of load balancing. Cisco routers will load-balance over a maximum of four equal-cost paths between a source and destination using either per-destination or per-packet load balancing when using OSPF. The default of per-destination is based on connectivity bandwidth at 56 Kbps or greater.

OSPF Convergence

Since OSPF is a link-state-based routing protocol, it adapts quickly to network topology changes. OSPF detects topology changes based on the interface status or the failure to receive a response to an OSPF HELLO packet of an attached neighbor within a given amount of time. OSPF has a default timer of 40 seconds in broadcast networks (such as LANs) and two minutes in non-broadcast networks (such as WANs).

The routes are recalculated by the router recognizing the failed link and sends a link-state packet to all the routers within the area. Each router then recalculates all the routes within its routing table.

OSPF Scalability

The addressing scheme, number of areas, and number of links within the OSPF network all affect the scalability of an OSPF network. Routers use memory for storing all the link states for each area to which a router belongs. The more areas attached to a router, the larger the table. Scaling OSPF therefore depends on the effective use of route summarization and stub areas to reduce memory requirements. The larger the link-state database, the more CPU cycles are required during the recalculation of the shortest-path-first algorithm.

Minimizing the size of an OSPF area and the number of links within the area, along with route summarization, enables OSPF to scale to large networks. OSPF only sends small HELLO packets and link-state updates when a topology change occurs or at startup. This is a great benefit for preserving bandwidth utilization, as compared to distance-vector routing protocols such as RIP or IGRP.


OSPF Security

OSPF can use an authentication field to verify that a router connecting as a neighbor is indeed a router that belongs within the network. OSPF routers by their very nature do not allow the filtering of routes, since all OSPF routers must have the same routing information within an area. Using authentication, an OSPF router can verify that it should exchange topology information with a new router that has joined the network. In this way, not only does OSPF provide some protection from unwanted access, it assists in keeping a stable network.

CHAPTER

5

Frame Relay Network Design



Frame Relay is based on a packet-switched data network. The differential of Frame Relay to previous packet-switched networks like X.25 is that Frame Relay switches a frame versus a packet. Frame Relay has considerable low overhead and its speed through the network is in part due to not ensuring the delivery of data. Frame Relay as a WAN network solution has grown because of its low cost for acceptable performance, as compared to leased-line WAN solutions. An optimal Frame Relay network design is based on the following:

- Balancing the cost savings of using a public network with the business performance requirements
- A scalable WAN design founded in a manageable environment
- Utilizes a hierarchical design

Main concerns for implementing a Frame Relay design is the ability of the design to scale to not only topology growth but to traffic growth. Components for creating a scalable Frame Relay network designs are as follows:

- The adherence to the three-layer router model of core, distribution, and access layers
- Overall hierarchical design
- Implementing various mesh topology design
- Addressing protocol broadcast issues
- Addressing performance concerns


Meeting these guidelines results in providing a scalable, high-availability, and low-cost Frame Relay network design.

Hierarchical Design of Frame Relay Internetworks

Frame Relay design is based on *permanent virtual connections* (PVCs). A PVC is identified using a *Data Connection Link Identifier* (DLCI) number. Multiple PVCs are possible over a single physical communication link. Using this capability, a single link can communicate with multiple locations. This function is shown in Figure 5-1 where router R1 using two PVCs communicates with two other routers over the public Frame Relay network.

Figure 5-1

Frame Relay PVCs connecting a single router to two routers



A PVC can be assigned a bandwidth. The total bandwidth of all defined PVCs can equal the actual bandwidth of the physical communication link. In a sense, Frame Relay acts as a *time-division multiplexer* (TDM) over a public network.

Due to the nature of Frame Relay services through PVCs, hierarchical designs are more logical than physical in definition. Each PVC may be guaranteed two bandwidth parameters, the *committed information rate* (CIR) and *excessive burst limits* (Be). The CIR is an agreement with the Frame Relay provider for a minimum throughput for the PVC. The excessive burst limit is an agreement with the Frame Relay provider for the available bandwidth for use by the PVC over and above the PVC bandwidth

to the maximum available on the physical link. These two variables greatly influence the cost and therefore the design of the Frame Relay network.

Frame Relay Scalability

Scalability is achieved in Frame Relay network design through the implementation of a hierarchy. Using a hierarchy enables incremental growth. The hierarchical approach, however, must follow the three-layer routing model in order for meeting high-availability, acceptable performance, and low-cost requirements. These requirements can be met through careful planning of actual performance requirements at remote locations, the degree of high-availability service, and minimizing the complexity of the hierarchy.

Frame Relay Management

Managing a hierarchical network is minimized through the partitioning of the network into smaller elements. By simplifying the network into manageable modules, troubleshooting is eased. The partitioning also provides protection against broadcast storms and routing loops. A hierarchical design inherently provides a flexible network topology, allowing the inclusion of other technologies into the network design. This leads to a hybrid approach for the overall network infrastructure. Although hybrid network design may enable greater service, it does make network management a bit more complex. Finally, router management in hierarchical Frame Relay networks is reduced due to fewer network connections based on the hierarchy.

Frame Relay Performance

Hierarchical network design lends itself to protecting networks from broadcast and multicast traffic issues. A regional hierarchy with smaller areas enables the Frame Relay network to maintain overall network performance requirements. Limiting the number of routers within an area or layer minimizes the chances of traffic bottlenecks due to broadcast traffic.

Frame Relay Network Topology

The network topology design chosen for implementing Frame Relay networks is dependent on many variables. Among these are the types of protocols supported and the actual traffic characteristics and patterns generated by applications using the network. It is recommended that an optimal Frame Relay network design support anywhere from a maximum of 10 to 50 PVCs per physical interface. Consider the following factors in determining the number of PVCs to support:


- Broadcast-intensive protocols constrain the number of PVCs. Segregating the protocols into their own PVC for better management requires more PVCs in multiprotocol networks.
- Broadcast updates due to routing protocols may consume bandwidth. The number, type, and frequency of the routing protocol updates will dictate the number of PVCs required to meet service levels.
- The available bandwidth of the physical Frame Relay connection as measured against the amount of broadcast traffic may dictate higher-bandwidth PVCs with higher CIRs and excess burst limits. Because each PVC has more bandwidth, however, the number of PVCs is reduced.
- Static routes can either eliminate or reduce the amount of broadcasts, thereby enabling more PVCs per physical connection.
- Large networks tend to create large routing protocol updates. Large updates and frequencies require higher bandwidth, thereby reducing the number of available PVCs per physical link.

The topology of a Frame Relay network is comprised of different design formats. Each format has its advantageous and disadvantageous. The network requirements along with the considerations outlined above on the number of PVCs required in a design need to be addressed in using the various topology layouts.

Frame Relay Star Topology

A Frame Relay star topology is depicted in Figure 5-2. The configuration is referred to as a star due to the single connection by all remote sites to a central location. Star topologies minimize the number of PVCs and result in a low-cost design. Due to its design, however, bandwidth at the central site

Figure 5-2
Frame Relay
star topology



becomes an issue since it becomes limited due to the number of remote locations connecting over the physical connection. Likewise, high-availability through alternate paths and rerouting of data from the remote locations is non-existent since there is only one path from the remote location to the rest of the network.

An advantage to a star topology is its ease of management, but its disadvantages make it a poor choice for basing a foundation on its network design. These drawbacks include the core or hub router as a single point of failure, performance problems of the backbone due to the single core router connection, and the inability of a star topology to scale.


Frame Relay Fully Meshed Topology

A fully meshed Frame Relay network provides a very high degree of availability. As shown in Figure 5-3, a fully meshed network uses PVCs connecting all Frame Relay points on the network. Disadvantageous to using a fully meshed network is the number of PVCs required. A PVC is required

for logically connecting to each router on the network. A fully meshed topology requires $[n(n-1)]/2$ PVCs, where n is the number of routers being connected to the Frame Relay network. For example, a fully meshed network of five routers requires $[5(5-1)]/2$, which equals 10 PVCs.

Although Frame Relay networks are *non-broadcast multi-access* (NBMA) networks, a router sends a broadcast over each active PVC. This replication process leads to excessive CPU and bandwidth requirements for routing updates, spanning tree updates and SAP updates.

Figure 5-3
Frame Relay fully
meshed topology



In small Frame Relay networks, a fully meshed topology is a reasonable design. The issues that make a fully meshed network for large networks a poor design are as follows:

- A large number of PVCs
- CPU and bandwidth overhead due to packet and broadcast replication
- Management complexity

Frame Relay Partially Meshed Topology

Merging the ease of design and management using a star topology with the high-availability feature provided by a fully meshed topology results in a requirements-balanced partially meshed topology. Seen in Figure 5-4, a partially meshed topology consists of two star topologies being supported by remote locations. Partially meshed topologies are ideal for regional implementation. Their advantages are as follows:

- High-availability
- Relatively low cost, as compared to fully meshed
- Minimum number of PVCs required
- Acceptable performance at a reasonable cost

Data must flow through one of the core routers for communication between locations of a partially meshed topology without a direct PVC.


Frame Relay Fully Meshed Hierarchical Topology

Applying the fully meshed topology to an overall hierarchy for the three layers of the routing layer model results in a design that scales and localizes traffic due to the creation of manageable segments. The modularity of the design enables the network as a whole to scale well. As shown in Figure 5-5, the hierarchy is based on the strategic connections made across the routing layer model.

Although, again, this topology provides high redundancy and modularity, it continues to have the packet-broadcast replication problem. The balance of service to cost is also lost due to the extra number of routers, physical links, and PVCs required.

Figure 5-4

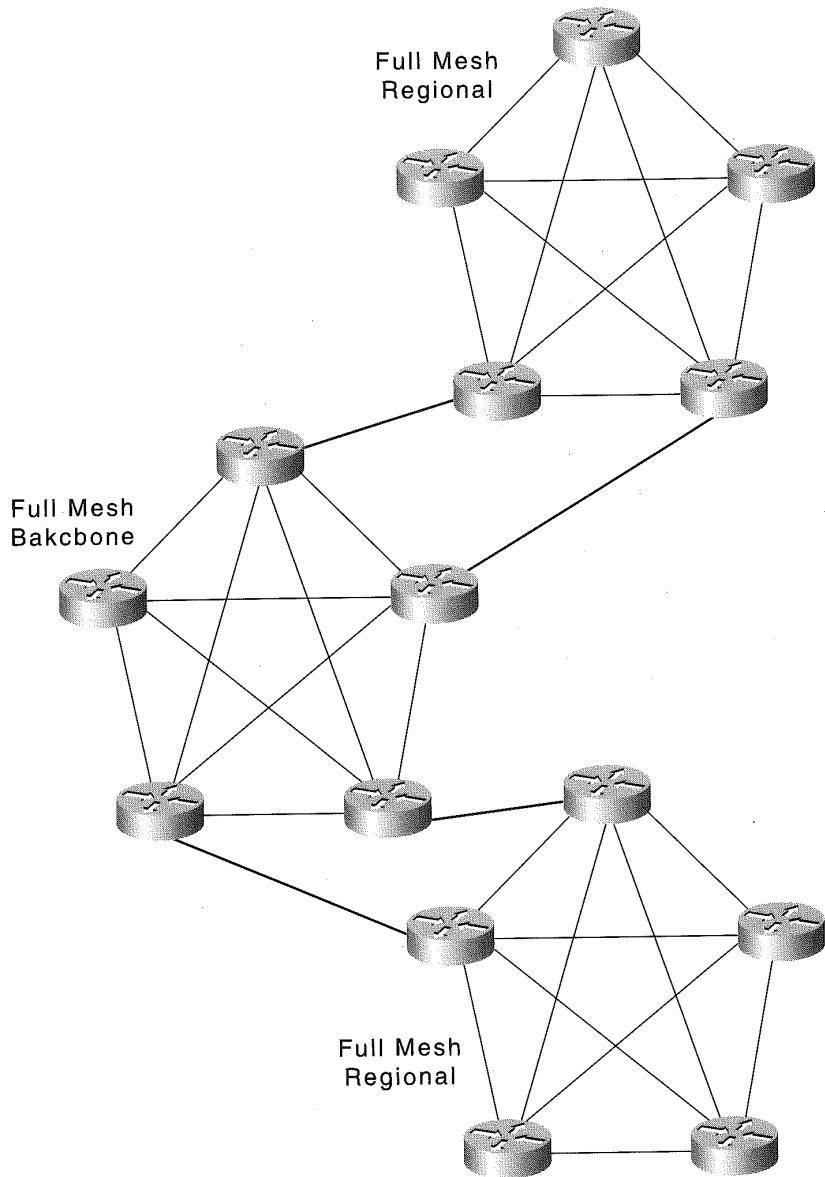
Frame Relay partial mesh topology



Frame Relay Hybrid Meshed Hierarchical Topology

Managing the balance between core backbone performance and maintaining a low-cost network design results in a hybrid hierarchical Frame Relay network. A hybrid hierarchical network, as depicted in Figure 5-6, uses private leased lines for creating a fully meshed backbone and partially or fully meshed Frame Relay networks for connection to the regional network.

Figure 5-5
Frame Relay fully
meshed hierarchical
topology



In Figure 5-6, we see the use of an ATM core backbone feeding a leased line distribution network. The distribution layer then provides network connectivity using a partially meshed topology. This topology provides high-availability, great bandwidth for the backbone, network segmentation, and simplified router configuration management.

Broadcast Traffic Issues

Broadcasts are typically used for routing protocols to update network devices on selecting the best path between two destinations on the network. Many routing protocols update their neighbors or peers on a periodic basis. Routers replicate a broadcast to every active PVC defined on the router for transmission to the partner node at the other end of the PVC. Figure 5-7 illustrates this point.

In managing the broadcasts of routing protocols, it is important to understand the time requirement for topology changes. In stable networks, the timers that manage the broadcast updates for individual routing protocols can be extended, which helps router and bandwidth overhead in supporting the routing protocol updates.

Another alternative is to include efficient routing protocols, such as EIGRP, in the design in order to reduce the routing protocol broadcast updates over the Frame Relay network. Managing the replication of broadcasts and packets is of paramount concern. Fully meshed networks actually increase the overall cost of a network and increase the overall load on the network. Table 5-1 lists the relative traffic levels as they relate to broadcast traffic generated by routing protocols.

Performance Considerations

Several factors affect the performance of Frame Relay networks. We have already discussed the effect of broadcasts on the network. Broadcasts are the primary concern for designing the bandwidth and number of PVCs necessary for a viable Frame Relay network. During the planning stage of

Figure 5-6

Use of private leased lines and regional Frame Relay networks




Figure 5-7
Broadcast replication
over Frame Relay
PVCs




Table 5-1

Relative Broadcast
Traffic Generated
by Popular Routing
Protocols

| Network Protocol | Routing Protocol | Relative Broadcast Traffic Level |
|--|--|----------------------------------|
| AppleTalk | Routing Table Maintenance Protocol (RTMP) | High |
| | Enhanced Interior Gateway Routing Protocol (EIGRP) | Low |
| Novell Internetwork Packet Exchange (IPX) | Routing Information Protocol (RIP) | High |
| | Service Advertisement Protocol (SAP) | High |
| | Enhanced Interior Gateway Routing Protocol (EIGRP) | Low |
| Internet Protocol (IP) | Routing Information Protocol (RIP) | High |
| | Interior Gateway Protocol (IGRP) | High |
| | Open Shortest Path First (OSPF) | Low |
| | Intermediate System -Intermediate System (IS-IS) | Low |
| | Enhanced Interior Gateway Protocol (EIGRP) | Low |
| | Border Gateway Protocol (BGP) | None |
| | Exterior Gateway Protocol (EGP) | None |
| DECnet Phase IV | DECnet Routing | High |
| DECnet Phase V | IS-IS | Low |
| International Organization for Standardization (ISO) Connectionless Network Service (CLNS) | IS-IS | Low |
| | ISO-IGRP | High |
| Xerox Network Systems (XNS) | RIP | High |
| Banyan Virtual Integrated Network Service (VINES) | Routing Table Protocol (RTP) | High |
| | Sequenced RTP | Low |

developing the Frame Relay network design, the following must be considered:

- Maximum rate requirements
- CIR
- Management of multiprotocol traffic

Determining Maximum rate

The Frame Relay provider uses several metrics to determine the billing of the Frame Relay connections. Therefore, it is important to fully understand the bandwidth and number of PVCs required to meet business service levels. The metrics used for determining the Frame Relay network configuration are

Committed burst (Bc): The number of bits committed to accept and transmit at the CIR

Excess burst (Be): The number of bits to transmit after reaching the Bc value

Committed Information Rate (CIR): The maximum permitted traffic level for each PVC

Maximum data rate (MaxR): Calculated value measured in bits per second; $(Bc + Be)/Bc * CIR$

Determination of the CIR, Bc, and Be is predicated on the actual speed of the physical line. The maximum values cannot extend past the maximum speed of the link. In addition, the application profiles will influence the metrics based on the type of service, transport mechanisms, and usage of each application using the PVCs.

Committed Information Rate (CIR)

The CIR is the guaranteed bandwidth the Frame Relay service provides for each PVC on the physical link. For example, a CIR of 19.2 Kbps on a 128 Kbps physical link commits the Frame Relay network to provide 19.2 Kbps of throughput for the PVC between source and destination. CIR is the metric most influential on the capability to meet the service levels for the applications. Failure to properly calculate the appropriate CIR level results in poor performance and failure to meet service levels.

Underestimating the CIR results in *discard-eligible* (DE) frames. The DE bit value is activated by a Frame Relay switch when the bandwidth used on the PVC begins to exceed the CIR. Frame Relay switches inspect the DE bit value within the frame. If the DE bit is on, the frame may be discarded based on the switches resource constraints, network congestion, and available bandwidth.

FECN/BECN Congestion Protocol

Frame Relay institutes a congestion protocol to protect network resources from overuse, known as FECN/BECN. *Forward Explicit Congestion Notification* (FECN) is a Frame Relay message used to notify a receiving device that there is a congestion problem. *Backward Explicit Congestion Notification* (BECN) is a Frame Relay message used to notify a sending device that there is a congestion problem. These messages enable the network devices to throttle the traffic onto the network. Cisco routers support the use of FECN and BECN.

Virtual Subinterface and Multiprotocol Management

Support for multiple protocols over Frame Relay connections requires some thought on traffic management. Cisco IOS enables the use of subinterfaces on physical interfaces. This capability to create virtual interfaces, diagrammed in Figure 5-8, enables a network designer to use all the tuning, reporting, and management functions of the Cisco IOS interface commands for each individual PVC.


Using this feature of virtual interfaces also creates unique buffers on the output queues for each PVC, versus one output buffer queue for the entire physical connection. The result is better performance and management using virtual subinterfaces.

SNA Support

Cisco IOS supports the transport of IBM *Systems Network Architecture* (SNA) protocols over Frame Relay using the RFC 1490/FRF.3 specification. The specification describes the encapsulation technique for transporting

Figure 5-8

The use of subinterfaces for dedicating protocols to PVCs



the SNA protocols. Cisco has applied their own algorithms for supporting enhanced features such as local acknowledgment, dynamic rerouting, SNA prioritization, and PVC prioritization.


Boundary Network Node (BNN)

Cisco routers implementing RFC 1490/FRF.3 can connect LAN-attached or SDLC-attached SNA resources directly to an IBM front-end processor without the use of a data-center-based router or any other intermediate Frame Relay device. The IBM front-end processor must be using *Network Control Program (NCP) V7.1* or higher *Boundary Network Node (BNN)* functions. Using a Cisco router at the remote location enables these SNA devices to maintain their current configuration while realizing the design benefits of a Frame Relay network. Figure 5-9 illustrates an SNA BNN connection to a mainframe front-end processor using Cisco routers at the remote location.

Locations having multiple SNA *physical units (PUs)* requiring connectivity may use a single PVC. This is accomplished by implementing a *Service Access Point (SAP)* multiplexing feature. Each SNA PU is assigned a unique SAP address, which enables the Cisco router to support multiple SNA PUs over the single PVC.

Figure 5-9

SNA BNN Frame Relay connectivity using Cisco routers




Boundary Access Node (BAN)

RFC1490/FRF.3 enhances Frame Relay connectivity directly to the FEP by including the IEEE 802.5 MAC header in every frame. This specification is called *Boundary Access Node* (BAN). Using BAN, an unlimited number of SNA devices are supported over a single Frame Relay PVC. BAN eliminates the need to use SAP addresses for multiplexing the SNA connections over a single Frame Relay PVC.

Additionally, BAN supports duplicate DLCI-MAC address mappings on the front-end processors for load balancing and redundancy. Support for BAN on the IBM front-end processor requires NCP V7.3 or higher, and the Cisco IOS must be using IOS 11.1 or greater. Figure 5-10 illustrates the use of BAN connectivity.

Figure 5-10
Cisco IOS and router support for BAN functionality



The differences between BNN and BAN are as follows:

- BAN does not greatly benefit reduced router configuration over BNN for single SNA PU connectivity.
- For LAN-attached SNA PUs, BNN requires a router configuration change, as opposed to the dynamic use of MAC addresses employed by BAN.
- BNN is more efficient for SDLC-attached devices than BAN. At locations that have both SDLC-attached and LAN-attached SNA PUs, a combination of BNN and BAN is beneficial.
- BAN may require an NCP upgrade to V7.3.
- Only BAN supports load balancing and dynamic redundancy.

FRAS Host Support

Cisco IOS supports the RFC 1490/FRF.3 node function at the data center router using the *Frame Relay access support* (FRAS) host function. As shown in Figure 5-11, instead of the Frame Relay PVC terminating at an IBM front-end processor, a Cisco router is used. The Cisco IOS SNA connectivity features for connecting to the mainframe using either SDLC, LAN, or a channel attachment with either a *channel interface processor* (CIP) or a *channel port adapter* (CPA) are then employed for completing the SNA connection.






Figure 5-11
Cisco FRAS Host support for BNN and BAN connectivity

CHAPTER

6

ATM
Internetworking
Design



Asynchronous Transfer Mode (ATM) is the first networking architecture developed specifically for supporting multiple services. ATM networks are capable of supporting audio (voice), video, and data simultaneously. ATM is currently able to support up to 2.5 Gbps of bandwidth. Data networks immediately get a performance enhancement when moving to ATM due to the increased memory. Voice networks realize a cost savings, due in part to sharing the same network with data and through voice compression, silence compression, repetitive pattern suppression, and dynamic bandwidth allocation. The ATM fixed-size 53-byte cell enables ATM to support the isochronicity of a *time-division multiplexed* (TDM) private network with the efficiencies of *public switched data networks* (PDSN).

Most network designers are first challenged by the integration of ATM with the data network. Data network integration requires legacy network protocols to traverse a cell-based switched network. ATM can accomplish this in several ways. The first of these is LAN emulation.

LAN Emulation (LANE)


ATM employs a standards-based specification for enabling the installed base of legacy LANs and the legacy network protocols used on these LANs to communicate over an ATM network. This standard is known as *LAN emulation* (LANE). LANE uses the *Media Access Control* (MAC) sublayer of the OSI data link control Layer 2. Using MAC encapsulation techniques enables ATM to address the majority of Layer 2 and Layer 3 networking protocols. ATM LANE logically extends the appearance of a LAN, thereby providing legacy protocols with equivalent performance characteristics, as are found in traditional LAN environments. Figure 6-1 illustrates a typical ATM topology with LANE support.

LANE uses ATM's *emulated LANs* (ELANs). Using ELANs, a LAN in one location is logically connected to a LAN in another location. This allows a network designer to extend a LAN over an ATM WAN, avoiding the need for routing packets between the two locations. LANE services can be employed by ATM-attached servers or workstations, edge devices such as switches, and routers when routing between ELANs is required.

ATM LANE uses four components to establish end-to-end connectivity for legacy protocols and devices:

- LAN Emulation Client
- *LAN Emulation Configuration Server* (LECS)

Figure 6-1
ATM LANE topology



- LAN Emulation Server (LES)
- Broadcast and Unknown Server (BUS)

LAN Emulation Client (LEC)

Any end system that connects using ATM requires a *LAN emulation Client (LEC)*. The LEC performs the emulation necessary in support of the legacy LAN. An LEC performs the following functions:

- Data forwarding
- Address resolution
- Registering MAC addresses with the LANE server

- Communication with other LECs using ATM *virtual channel connections* (VCCs)
End systems that support the LEC functions act as
 - ATM-attached workstations
 - ATM-attached servers
 - ATM LAN switches (Cisco Catalyst family)
 - ATM attached routers (Cisco 12000, 7500, 7000, 4700, 4500, and 4000 series)

LAN Emulation Configuration Server (LECS)

The ELAN database is maintained by the *LAN emulation configuration server* (LECS) and is manually updated by the router administrator. In addition, the LECS builds and maintains an ATM address database of *LAN Emulation Servers* (LES). The LECS maps an ELAN name to a LES ATM address. The LECS performs the following LANE functions:

- It accepts queries from a LEC.
- It responds to a LEC query with an ATM address of the LES for the ELAN.
- It serves multiple ELANs.
- It is manually defined and maintained.

The LECS assigns individual clients to a ELAN by directing them to the LES that corresponds to the ELAN.

LAN Emulation Server (LES)

LECs are controlled from a central control point called a *LAN Emulation Server* (LES). LECs communicate with the LES using a *Control Direct Virtual Channel Connection* (VCC). The Control Direct VCC is used for forwarding registration and control information. The LES uses a *Control Distribute VCC*, a point-to-multipoint VCC, enabling the LES to forward control information to all the LECs. The LES services the *LAN Emulation Address Resolution Protocol* (LE_ARP) request, which it uses to build and maintain a list of LAN destination MAC addresses.

Broadcast Unknown Server (BUS)

ATM is based on the notion that the network is point-to-point. Therefore, there is no inherent support for broadcast or any-to-any services. LANE provides this type of support over ATM by centralizing broadcast and multicast functions on a *Broadcast And Unknown Server* (BUS). Each LEC communicates with the BUS using a Multicast Send VCC. The BUS communicates with all LECs using point-multipoint VCC, known as the Multicast Forward VCC. A BUS reassembles received cells on each Multicast Send VCC in sequence to create the complete frame. Once a frame is complete, it is then sent to all the LECs on a Multicast Forward VCC. This ensures the proper sequence of data between LECs.

LANE Design Considerations

The following are guidelines for designing LANE services on Cisco routers:

- The Cisco AIP has a bidirectional limit of 60 thousand packets per second (pps).
- The ATM interface on a Cisco router has the capability of supporting up to 255 subinterfaces.
- Only one active LECS can support all the ELANs. Other LECS operate in backup mode.
- Each ELAN has one LES/BUS pair and one or more LECs.
- LES and BUS must be defined on the same subinterface of the router AIP.
- Only one LES/BUS pair per ELAN is permitted.
- Only one active LES/BUS pair per subinterface is allowed.
- The LANE Phase 1 standard does not provide for LES/BUS redundancy.
- The LECS can reside on a different router than the LES/BUS pair.
- VCCs are supported over *switched virtual circuits* (SVCs) or *permanent virtual circuits* (PVCs).
- A subinterface supports only one LEC.
- Protocols such as AppleTalk, IP, and IPX are routable over a LEC if they are defined on the AIP subinterface.
- An ELAN should be in only one subnet for IP.

Network Support

The LANE support in Cisco IOS enables legacy LAN protocols to utilize ATM as the transport mechanism for inter-LAN communications. The following features highlight the Cisco IOS support for LANE:

- Support for Ethernet-emulated LANs only. There is currently no Token Ring LAN emulation support.
- Support for routing between ELANs using IP, IPX, or AppleTalk
- Support for bridging between ELANs
- Support for bridging between ELANs and LANs
- LANE server redundancy support through *simple server redundancy protocol* (SSRP)
- IP gateway redundancy support using *Hot Standby Routing Protocol* (HSRP)
- DECnet, Banyan VINES, and XNS-routed protocols

Addressing

LANE requires MAC addressing for every client. LANE clients defined on the same interface or subinterface automatically have the same MAC address. This MAC address is used as the *end system identifier* (ESI) value of the ATM address. Although the MAC address is duplicated, the resulting ATM address representing each LANE client is unique. All ATM addresses must be unique for proper ATM operations. Each LANE services component has an ATM address unique from all other ATM addresses.

LANE ATM Addresses

LANE uses the NSAP ATM address syntax. The address format used by LANE is as follows:

- A 13-byte prefix that includes the following fields defined by the ATM Forum:
 - An *Authority and Format Identifier* (AFI) field (one byte)
 - A *Data Country Code* (DCC) or *International Code Designator* (ICD) field (two bytes)

- A *Domain Specific Part Format Identifier* (DFI field) (one byte)
- An Administrative Authority field (three bytes)
- A Reserved field (two bytes)
- A Routing Domain field (two bytes)
- An Area field (two bytes)
- A six-byte *end-system identifier* (ESI)
- A one-byte selector field

Cisco's Method of Automatically Assigning ATM Addresses

The Cisco IOS supports an automated function of defining ATM and MAC addresses, which are used in the LECS database. The automation process uses a pool of eight MAC addresses that are assigned to each router ATM interface. The Cisco IOS applies the addresses to the LANE components using the following methodology:

- All LANE components on the router use the same prefix value, which identifies a switch and must be defined within the switch.
- The first address in the MAC address pool becomes the ESI field value for every LANE client on the interface.
- The second address in the MAC address pool becomes the ESI field value for every LANE server on the interface.
- The third address in the MAC address pool becomes the ESI field value for the LANE broadcast-and-unknown server on the interface.
- The fourth address in the MAC address pool becomes the ESI field value for the LANE configuration server on the interface.
- The selector field for the LANE configuration server is set to a 0 value. All other components use the subinterface number of interface to which they are defined as the selector field.

The requirement that the LANE components be defined on different subinterfaces of an ATM interface results in a unique ATM address, due to the use of the selector field value being set to the subinterface number.

Using ATM Address Templates

ATM address definitions are greatly simplified through the use of address templates. However, these templates are not supported for the E.164 ATM address format. The address templates used for LANE ATM addressing can use either an asterisk (*) or an ellipsis (...) character. An asterisk is used for matching any single character, while an ellipsis is used for matching leading or trailing characters. Table 6-1 lists the address template value determination.

The ATM address templates can be either a prefix or an ESI template. When using a prefix template, the first 13 bytes match the defined prefix for the switch but use wildcards for the ESI and selector fields. An ESI template matches the ESI field but uses wildcards for the prefix and selector fields.

Rules for Assigning Components to Interfaces and Subinterfaces

The LANE components can be assigned to the primary ATM interface as well as the subinterfaces. The following are guidelines for applying LANE components on a Cisco router ATM interface:

- The LECS always runs on the primary interface.
- Assigning a component to the primary interface falls through to the subinterface 0 definition.
- The LES and LEC of the same emulated LAN can be configured on the same subinterface in a router.

Table 6-1

Determining ATM Address Template Values

| Unspecified Digits In | Resulting Value Is |
|----------------------------|--|
| Prefix (First 13 bytes) | Obtained from ATM switch via Interim Local Management Interface (ILMI) |
| ESI (Next six bytes) | Filled using the first MAC address of the MAC address pool plus <ul style="list-style-type: none"> 0-LANE client 1-LANE server 2-LANE broadcast-and-unknown server 3-LANE configuration server |
| Selector field (last byte) | Subinterface number, in the range 0 through 255. |

- LECs of two different emulated LANs must be defined on a different subinterface in a router.
- LESs of two different emulated LANs must be defined on a different subinterface in a router.

Redundancy in LANE Environments

The ATM LANE V 1.0 specification does not provide for redundancy of the LANE components. High availability is always a goal for network designers and the single point of failure in the LANE specification requires a technique for redundancy. Cisco IOS supports LANE redundancy through the implementation of Simple Server Replication Protocol (SSRP).

SSRP supports the redundancy of LECS and LES/BUS services. LECS redundancy is provided by configuring multiple LECS addresses in the ATM switches. Each defined LECS is defined with a rank, which is the index (the number of the entry in the LECS address table) of the LECS address in the table.

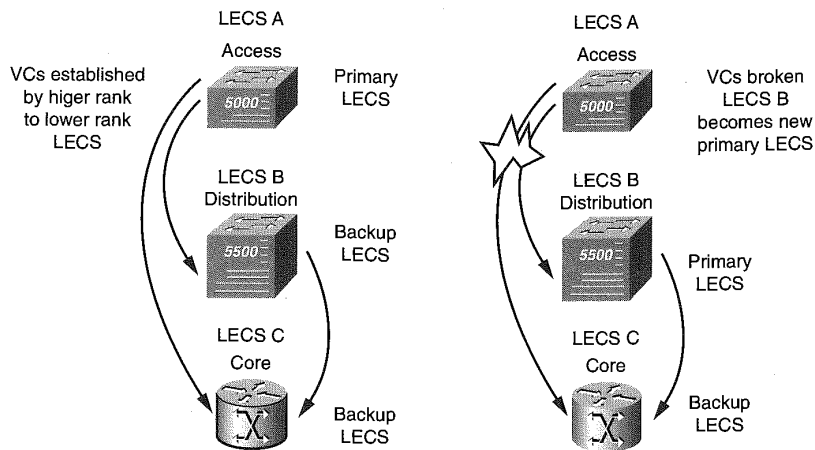
At initialization the LECS requests the LECS address table from the ATM switch. The requesting LECS on receipt of the LECS address table tries to connect to all the LECSs with a lower rank. In this way, the LECS learns of its role in the redundancy hierarchy. A LECS that connects with a LECS whose rank is higher places itself in a backup mode. The LECS that connects to all other LECS and does not find a ranking higher than its own assumes the responsibility of the primary LECS.

In this hierarchy, as shown in Figure 6-2, the failure of a primary LECS does not result in a LANE failure. Rather, the second-highest ranking LECS assumes the primary LECS role. Loss of the VCC between the primary and highest ranking secondary signals the highest secondary ranking LECS that it is now the primary LECS.

In theory, any number of LECS can be designed using SSRP. Cisco recommends, however, that no more than three LECS be designed into SSRP. This recommendation is based on adding a degree of complexity to the network design, which can lead to an increase in the time it takes for resolving problems.

LES/BUS redundancy using SSRP is similar in that it uses a primary-secondary hierarchy; however, the primary LES/BUS pair is assigned by the primary LECS. The LECS determines the primary LES/BUS pair by determining the LES/BUS pair having the highest priority with an open VCC to the primary LECS. The LES/BUS pair priority is assigned during configuration into the LECS database.

Figure 6-2
SSRP configuration
for LECS redundancy




The following guidelines are highly recommended for designing the LECS redundancy scheme and ensuring a properly running SSRP configuration:

- Each LECS must maintain the same ELAN database.
- Configure the LECS addresses in the LECS address table in the same order on each ATM switch in the network.
- Do not define two LECSs on the same ATM switch when using the Well Known Address. Only one of the LECS will register the Well Known Address with the switch, which can lead to initialization problems.

A second type of redundancy mechanism used in LANE is specific to ELANS using IP protocol. The *Host Standby Router Protocol* (HSRP), developed for traditional IP LAN topologies, enables two routers to share a common virtual IP address using a virtual MAC address assigned to the resulting virtual interface. This enables two routers to respond as the single IP gateway address for IP end stations. Figure 6-3 illustrates the use of HSRP with LANE.

The primary and secondary router interface is determined by the definition of HSRP on the interface or subinterface. HSRP exchanges definition information between the two routers to determine which interface is the primary gateway address. The secondary then sends HELLO messages to the primary to determine its viability. When the secondary does not receive a HELLO message from the primary HSRP router, it assumes the primary role.

Figure 6-3
HSRP configuration
for IP redundancy of
LANE clients



Data Exchange Interface (DXI)

ATM networks connect to serial-attached routers by implementing the ATM *data exchange interface* (DXI) specification. The DXI specification enables ATM *user-network interface* (UNI) connectivity between a Cisco router with only a serial interface to the ATM network. This is accomplished using an ATM *Data Service Unit* (ADSU).


As shown in Figure 6-4, router R1 connects to the ADSU using a *High Speed Serial Interface* (HSSI) connection. The ADSU receives data from the router in the ATM DXI format. The ADSU then converts the data into ATM cells and forwards them to the ATM network. The ADSU performs the opposite function for data going to the router.

Supported Modes

Although ATM DXI comes in three modes, the Cisco IOS supports only mode 1a. The three modes are

- Mode 1a, which supports AAL5 only, a 9232-octet maximum, and a 16-bit FCS, with up to 1,023 virtual circuits.
- Mode 1b, which supports AAL3/4 and AAL5, a 9224-octet maximum, and a 16-bit FCS. AAL5 supports up to 1,023 virtual circuits and AAL3/4 is supported on one virtual circuit.
- Mode 2, which supports AAL3/4 and AAL5 with 16,777,215 virtual circuits, a 65535-octet maximum, and a 32-bit FCS.

Figure 6-4
ATM DXI connectivity
for Cisco routers



DXI Addressing


DXI addressing uses a value that is equivalent to a Frame Relay data link connection identifier. In DXI, this field is called a DFA. The ADSU maps the DFA to the appropriate ATM *Virtual Path Identifier* (VPI) and *Virtual Connection Identifier* (VCI). Figure 6-5 illustrates the bytes and position mapping of the DXI DFA address to the ATM cell VPI and VCI values.

Classical IP

Cisco routers are configurable as both an IP client and an IP server in support of Classical IP. Classical IP enables the routers to view the ATM network as a *Logical IP Subnet* (LIS). Configuring the routers as an ATM ARP server enables Classical IP networks to communicate over an ATM network. The benefit to this is a simplified configuration. Classical IP support using an ATM ARP server alleviates the need to define the IP network address and ATM address of each end device connecting through the router in the router configuration.

ATM uses PVCs and SVCs and the ATM ARP server feature of Classical IP is SVC-specific. Using the ATM ARP server feature, each end device only configures its own ATM address and the address of the ATM ARP server. Since RFC 1577 allows for only one ATM ARP server address, no redundancy is available in Classical IP.

Figure 6-5
DXI address mapping
to ATM VPI/VCI
values



As shown in Figure 6-6, the ATM ARP server address can point to a Cisco router. IP clients using Classical IP make a connection to the ATM ARP server address defined in their configuration. The server then sends an *ATM Inverse ARP* (InARP) request to the client. The client responds with its IP network address and ATM address. The ATM ARP server places these addresses in its cache, which is used to resolve ATM ARP requests from IP clients. The IP client establishes a connection to the IP-ATM address provided in the ATM ARP server reply.

Multiprotocol over ATM (MPOA)


MPOA provides a single solution for transporting all protocols through an ATM network. MPOA V1.0, in concert with LANE *User-to-Network Interface* (UNI) V2.0, allows routers and other ATM networking devices to fully exploit VLANs, QoS, and high-availability. These network enhancements enable designers to add services while relieving traffic congestion and flexibility to the network. The key benefits to MPOA are as follows:

- Inter-VLAN “cut-through,” which maximizes bandwidth and network segmentation.
- Robust Layer 3 QoS features to support packetized traffic, such as video or voice, while ensuring data service levels.
- A software-only upgrade, which minimizes the cost and simplifies implementation.

The MPOA specification is built on four components:

- *MPOA Client* (MPC)
- *MPOA Server* (MPS)
- *Next Hop Resolution Protocol* (NHRP)
- *LAN Emulation* (LANE)

Figure 6-6
Classical IP support
on Cisco routers



Both MPC and MPS functions are supported on Cisco routers. MPOA uses a direct *virtual channel connection* (VCC) between the ingress (inbound) and egress (outbound) edge or host device. Direct VCCs are also termed shortcut VCCs. The direct VCC enables the forwarding of Layer 3 packets, normally routed through intermediate routers, between the source and destination host, thereby increasing performance and reducing latency.


Figure 6-7 illustrates the use of MCP, MPS, and NHRP for establishing a direct VCC between two edge devices servicing two end stations.

Multiprotocol Client (MPC)

Typically, the *Multiprotocol client* (MPC) will reside on an ATM edge device, such as a Cisco Catalyst family of switches. However, a Cisco router can perform the functions of an MPC or MPS. An MPC provides the following functions:

- Ingress/egress cache management
- ATM data-plane and control-plane VCC management

Figure 6-7
The MPOA flow for establishing a direct VCC between end stations



- MPOA frame-processing
- MPOA protocol and flow detection
- Identification of packets sent to an MPOA-capable router
- Establishes a direct VCC with the egress MPC

Multiprotocol Server (MPS)

The *Multiprotocol server* (MPS) provides the forwarding information used by the MPCs. The MPS maintains the information by using *Next Hop Resolution Protocol* (NHRP). MPS interacts with the NHRP module running in the router. MPS interacts with NHRP in the following manner:

1. The MPS converts the MPOA resolution request to a NHRP request. The MPS then sends the NHRP request to either the Next Hop MPS or the *Next Hop Server* (NHS) based on the results from the next hop information search through the MPS tables. MPS ensures that the correct encapsulation is used depending on the NHS type.
2. If the next hop is determined to be on a LANE cloud, the NHS sends resolution requests to the MPS. Likewise, the NHS sends resolution requests when the destination of the packet is unknown. The MPS can also request the NHS to terminate the request or discard the packet.
3. If the replies terminate in the router or the next hop interface uses LANE, resolution replies are sent from the NHS to the MPS.
4. Upon receiving resolution replies from the NHS, the MPS sends a MPOA resolution reply to the MPC.

MPS uses a network ID. The default network ID for all MPSs is one. Using different network IDs allows the network designer to segregate traffic. This enables the designer to permit direct VCCs between groups of LECs and deny direct VCCs between others. The network ID of an MPS and NHRP on the same router must be the same in order for requests, replies, and shortcuts to be transmitted across the MPS and NHRP.

MPOA Guidelines

The following is a list of guidelines for designing MPOA:

- An ELAN identifier must be defined for each ELAN.
- An MPC/MPS can serve as a single LEC or multiple LECs.
- A LEC can associate with any MPC/MPS.
- A LEC can attach to only one MPC and one MPS at a time.
- A LEC must break its attachment to the current MPC or MPS before attaching another MPC or MPS.
- A primary ATM interface can have multiple MPCs or MPSs defined with different control ATM addresses.
- Multiple MPCs or MPSs can be attached to the same interface.
- The interface attached to the MPC or MPS must be reachable through the ATM network by all LECs that bind to it.

Bandwidth Support on Routers

ATM is supported on the Cisco 7500 and 7000 series routers using the *ATM Interface Processor* (AIP). In designing the ATM internetwork in support of LANE, the total ATM bandwidth support for the entire router should not exceed 200 Mbps in full-duplex mode. This results in the following possible hardware configurations:

- Two *Transparent Asynchronous Transmitter / Receiver Interface* (TAXI) connections
- An OC-3 *Synchronous Optical Network* (SONET) and a E3 connection
- An OC-3 SONET and a low-use OC-3 SONET connection
- Five E3 connections

Configurable Traffic Parameters


The AIP provides the capability to shape various traffic. It supports up to eight rate queues. Each queue is programmed for a different peak rate. The ATM virtual circuits can be assigned to one of the eight rate queues. A virtual circuit can have an average rate and a burst size defined. The AIP supports the following configurable traffic rate parameters:

- Forward peak cell rate
- Backward peak cell rate
- Forward sustainable cell rate
- Backward sustainable cell rate
- Forward maximum burst
- Backward maximum burst

CHAPTER

7

Switched LAN Design



Local area networks (LANs) were initially developed to enable communities to share computing resources. As LANs grew in popularity, it became evident that the sharing of resources also mandated higher performance requirements. Higher performance is gained by dedicating what was once a shared medium used by many resources to a single network device. The dedication of a shared medium to a single device required the development of a networking hub that could itself become the shared medium, enabling the network devices to continue to share networking resources. LAN switching is the resulting technology.

As shown in Figure 7-1, a shared Ethernet hub connects multiple users to a router. The hub emulates the broadcast topology of an Ethernet and sends each Ethernet frame to all connected resources. In a switching environment, only Ethernet frames destined for the attached resource flow on the media. This is accomplished by the LAN switch, which uses OSI Layer 2 addressing. The *medium access control* (MAC) address is used for making the decision. In essence, the switch is a high-speed multiport bridge. This type of switching is called Layer 2 switching.

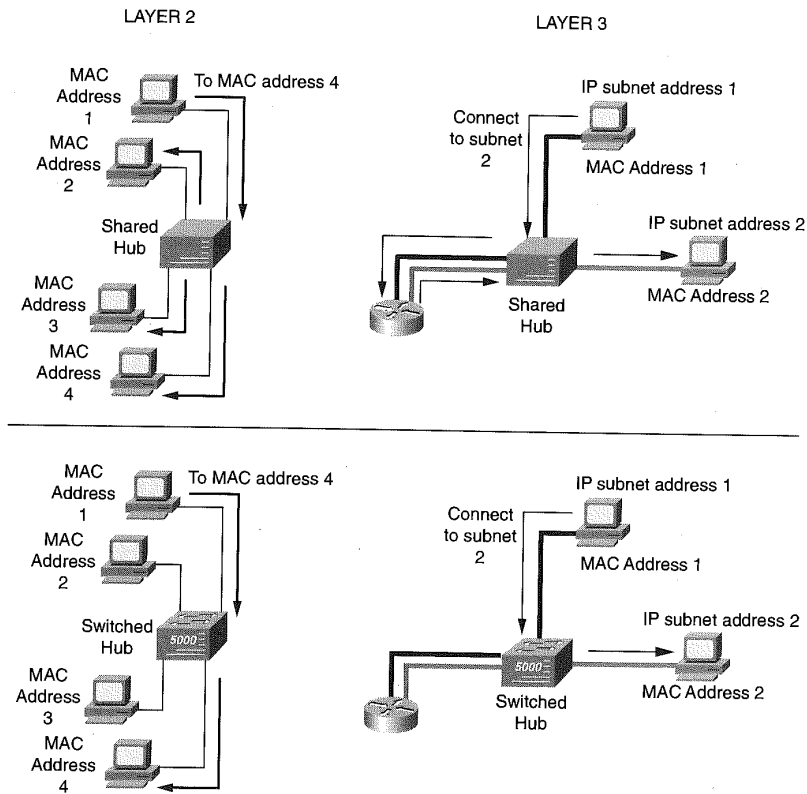
Enhancements to the switching technology paradigm have moved up the OSI Reference Model to the network layer. The network layer, Layer 3, supports routing protocols such as IP, Novell IPX, and AppleTalk. The switching intelligence now includes the capability to direct frames based on the Layer 3 addressing, versus the Layer 2 addressing.

In a shared hub environment, all devices attached to the shared media connection of the hub must be on the same IP subnet or network. A switched LAN environment enables multiple IP subnets or networks to coexist on the same switching medium. Communication between the IP subnets or networks is accomplished using Layer 3 switching, which is, in essence, the function of a router.

Switched LAN Factors

Many of the same issues that pertain to router networks are found in the switching networks. Because switching networks are based on the concept of bridging broadcast storms, offnet traffic and administration become factors in designing a viable switched network topology.

Figure 7-1
Shared LAN compared to switched LAN



Broadcast Radiation

Switching, though faster than routing, must still address the issue of broadcast storms or broadcast radiation. As is the theme with maintaining peak performance on the network, switching designs must minimize the affect of broadcasts. The use of *virtual LANs* (VLANs) in a switched network design helps to minimize the affect of broadcasts, but this results in scaling the number of hosts that can be supported on a VLAN. Using routers, a VLAN can be subsegmented based on traffic patterns, enabling a scalable network design. In such a design, multiple routers are required to avoid overutilization, due to the amount of VLAN traffic serviced by the router.

Well-behaved VLANs

VLANs were introduced at a time when the 80/20 rule was still in effect. The term “well-behaved VLAN” is used to characterize a VLAN that has 80 percent or more of its traffic local to the VLAN. Designing a “well-behaved VLAN” requires segmentation of services as well as devices. That is, client-server application database servers, e-mail servers, and file and print servers should all be dedicated to each VLAN. Although this type of design bodes well for network performance, management, and troubleshooting, it causes a strain on the network budgeting process. Since the term “well-behaved VLAN” was coined, we have seen a reversal of the 80/20 rule, where now clearly 80 percent of the VLAN traffic is destined outside the VLAN due to service consolidation.

Inter-VLAN Available Bandwidth

Routers are required to transport traffic between VLANs. A switched LAN design must account for the new 20/80 paradigm of traffic going off the VLAN. Therefore, the bandwidth connecting the various networking resources between the source and destination VLAN must be enough to handle the traffic load to meet service-level requirements.

Administrative Boundaries

In defining the switched network design, administrative boundaries must be given consideration. This is because switching tends to have a flattening affect on the operation of a network. Switching outside of the administrative boundary may affect the network within the administrative boundary, thereby causing poor performance or possibly disruption of service.

Cisco VLAN Implementation Support

VLANs are a logical grouping of devices. Typically, these devices are grouped either by function, department, floor, or for segmenting the LAN. Figure 7-2 illustrates a typical VLAN implementation and identifies six workstations on different ports of different switches on different floors.

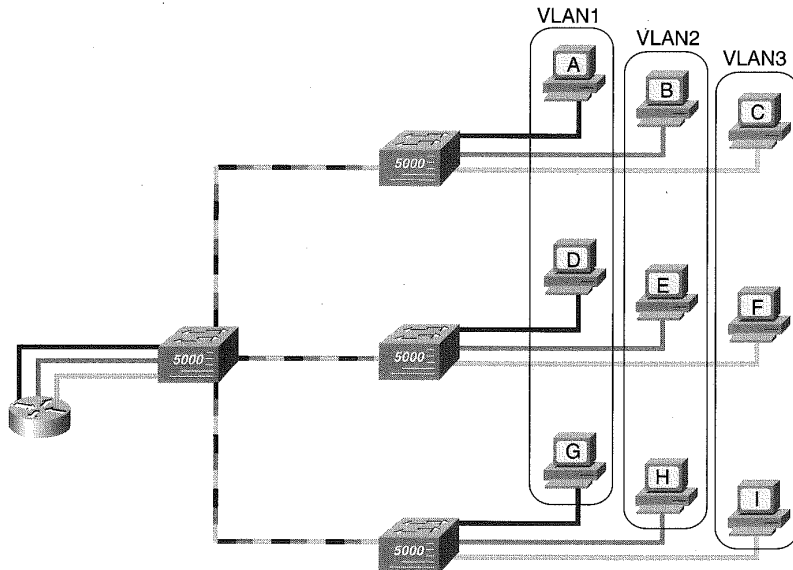
VLAN1 logically defines a “broadcast domain” on the switches with Stations A, D, and G. Likewise, the other stations should follow suit. The router in Figure 7-2 connects to all the VLANs, enabling VLAN-VLAN communication.

The first inception of VLANs is based on the OSI Layer 2 bridging and multiplexing mechanisms. The design of one host per port provides both high bandwidth for the attached host and ease of VLAN configuration.

Since switching relegates a network topology to a flat network, scalability and management of broadcast storms is made possible through the implementation of VLANs and routers. Combining the two realizes the following advantages:

- VLANs further refine the isolation of collision domains provided by a switching environment by only forwarding traffic within the VLAN. This contains broadcast and multicast traffic within the bridging domain created by the VLAN.
- VLANs enhance security since the logical groupings disallow Layer 2 inter-VLAN communication and require a router for Layer 3 VLAN communication. Inter-VLAN communication through the router allows a network engineer to employ the security features of the Cisco router by filtering appropriate packets from reaching other VLANs.

Figure 7-2
VLAN topology



- Logically grouping high-bandwidth end users together on a VLAN provides improved performance for devices requiring less bandwidth. This segregation of users enables performance design to meet service levels.
- The logical grouping of users on VLANs enables the moves, additions, and changes that frequently occur on a network to take place almost dynamically, rather than waiting for the altering of the actual infrastructure.

There are three types of methods for grouping devices on switches when defining VLANs:

- *Port-based VLANs:* Logically grouping VLANs based on the physical ports is known as segment-based VLANs. Only one VLAN is defined for a port and multiple VLANs are available for assignment on a switch. Intra-VLAN traffic within a switching hub is switched and must first be passed to a router, which then routes the traffic to the destined VLAN. Since port-based VLANs do not recognize Layer 3 addressing, all VLANs supporting Layer 3 traffic, such as IP, IPX, or AppleTalk, must be defined on the same network within the same VLAN.
- *Protocol-based VLANs:* The use of Layer 3 addressing enables a VLAN to differentiate between different network protocols. This type of VLAN is called a virtual subnet VLAN. Using network Layer 3 addressing as the scheme for defining VLANs enables more than one VLAN per port.
- *User-defined values:* This type of VLAN enables the network engineer to design a VLAN based on any field value within a packet. This provides the most flexibility along with granularity in defining VLANs. VLANs can be defined based on the type of service being requested or advertised.

Cisco VLAN implementation supports the following standards:

- IEEE 802.10
- IEEE 802.1d
- ATM LANE
- Cisco's own *Inter-Switch Link (ISL)*

These functions enable the design to handle multiple, disjointed, and overlaid broadcast groups on a flat network. Cisco currently employs port-based VLANs.

IEEE 802.10

The IEEE 802.10 standard defines two strategies for bridging the VLANs over a *metropolitan area network* (MAN) or *wide area network* (WAN) backbone in support of intra-VLAN communications. These strategies provide VLAN support for switching and routing.

The IEEE 802.10 standard implements a VLAN ID between the source MAC address and the *Link Service Access Point* (LSAP) fields of an FDDI frame prior to leaving the switch. The VLAN ID is a four-byte value within this header. The receiving switch examines the header to determine if the VLAN ID exists on the switch. If it does, it removes the IEEE 802.10 header and forwards the frame to the interfaces that match the VLAN ID. Only one VLAN ID is allowed per end user interface. The interface connecting to the MAN or WAN FDDI network is considered a VLAN trunk and supports multiple VLAN IDs.

In Figure 7-3, we see three switches connected over a MAN FDDI backbone. VLAN1 is defined on switch A, B, and C. The VLAN ID is 10. Switch A and B has VLAN2 with VLAN ID 20. Switch A and C has VLAN3 with VLAN ID 30. Using IEEE 802.10, the switches are able to forward packets between them, connecting the virtual LANs over the MAN or WAN.

In a routing scenario, IEEE 802.10 enables a split network by “gluing” the network together using a bridged path between the routers. This must include the enforcement of a same network addressing scheme for the VLANs being glued. Figure 7-4 illustrates the IEEE 802.10 routing configuration.

Some design considerations for using IEEE 802.10 include the following:

- Cisco routers use fast-switching mechanisms with IEEE 802.10.
- VLANs should be designed without splitting between locations over the backbone:
 - If subnets are split, they need to be “glued” by a bridged path.
 - VLANs must adhere to normal routing behavior in order for inter-VLAN communications to take place.

Table 7-1 lists the pros and cons of IEEE 802.10 switching and routing.

Figure 7-3
IEEE 802.10
connections for a
switched VLAN
environment

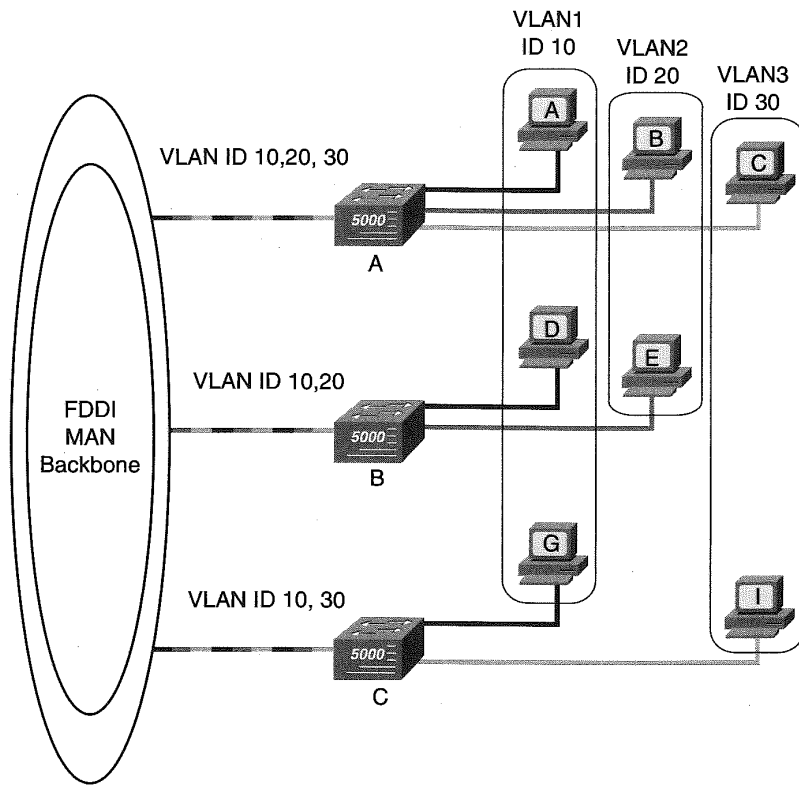



Table 7-1
Advantages and
Disadvantages of
Switching and
Routing when
Employing IEEE
802.10 VLANs

| Switched Backbone | | Routed Backbone | |
|--|--|--|--|
| Advantages | Disadvantages | Advantages | Disadvantages |
| Propagates VLAN ID information across the network | Backbone is running bridging. | No bridging in backbone | VLAN ID information is not propagated across backbone and must be configured manually. |
| Extends bridging domains, thereby enabling greater scaling | Broadcast traffic is high on the backbone. | Easily integrates into an existing network | Split subnets communicate using a bridged path between switches. |
| | | Can run native protocols in the backbone | |

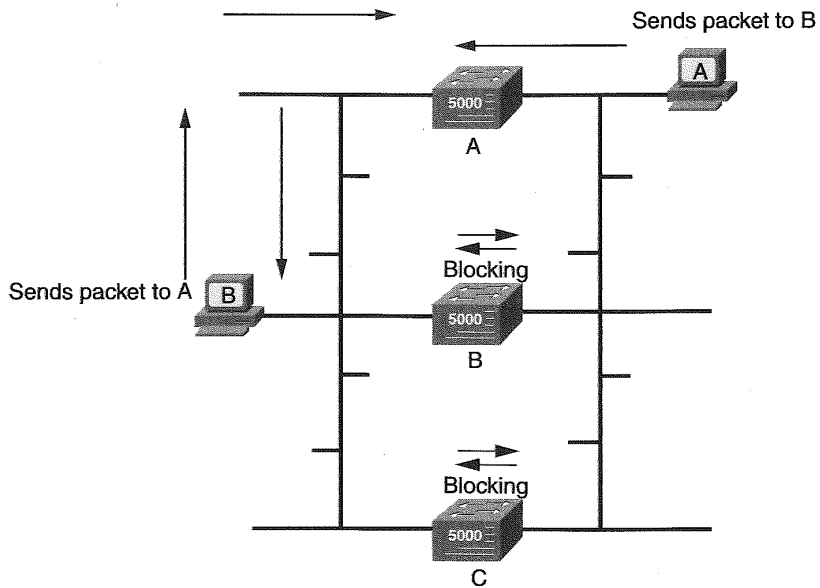
Figure 7-4
IEEE 802.1Q
connections for a
routed VLAN
environment



IEEE 802.1d

In high-availability designs, such as that shown in Figure 7-5, loops are possible, due to the replication of packets. As shown in the figure, switch A and B provide redundancy for a VLAN segment. Station A sends a packet to station B. Due to the high-availability design, the packet is looped back to segment A from switch B. The IEEE 802.1d standard prevents looping by employing the spanning tree protocol and placing ports in blocking mode, allowing a loop-free topology. Routers inherently prevent loops by using optimal paths.

Figure 7-5
IEEE 802.1d
spanning-tree
protocol for
preventing
bridging loops




Inter-Switch Link (ISL)

Cisco has taken the IEEE 802.10 specification and applied it to 100-Mbps Ethernet links between switches and routers. This proprietary specification is called the *Inter-Switch Link* (ISL) protocol. Shown in Figure 7-6, 100-Mbps Fast Ethernet links connect switch A to switch B, and then switch B to router C and switch D. Station A on switch A sends a packet to station D on switch D. Switch A determines the destination as being destined for a different switch for the same VLAN. ISL sends a 30-byte header to the Ethernet frame, which contains a two-byte VLAN ID field, prior to sending the packet over the Fast Ethernet link to switch B. Switch B interprets the ISL header and forwards the packet to the router C, which analyzes the ISL header and forwards the packet to switch D. At this point, switch D determines that the packet is for a station on a VLAN defined to it and forwards the packet over the appropriate port.

Switched LAN Design

Figure 7-6
Cisco ISL for VLAN,
spanning over Fast
Ethernet connections



LAN Emulation

Using ATM LANE specifications, VLANs can be mapped to an ELAN, enabling switches and routers to handle the majority of Layer 3 protocols over VLANs. Routers are used for inter-ELAN communications, which map to inter-VLAN communications. An ATM switch handles intra-ELAN communications.

Virtual Multi-homed Servers

Every network supports servers for corporate-wide access that are typically Web or e-mail servers. The mapping of VLANs to ATM ELANs enables multi-homed servers, thereby eliminating the need for routing packets to the server. ATM NICs allow up to eight ELAN and VLAN definitions. The


advantage here is that all packets of a VLAN defined to the end user domain and the server pass from switch to switch without the use of the router. The disadvantage is that the server ends up receiving all the broadcasts and multicasts of each ELAN defined to the NIC of the server. This may cause excessive overhead on the server, outside of providing services to the end user domain.

Figure 7-7 illustrates a multi-homed ATM server configuration. For non-ATM network topologies, the server must support multiple NICs for connecting to multiple VLANs or NICs supporting the backbone VLAN-trunking technology, IEEE 802.10 or ISL.

Switched LAN Topologies

Switched LAN topologies have quickly become the current network topology of choice. The switching topology designs fall into three generic categories:

Figure 7-7
Multi-home server
configuration using
ATM LANE topology



- Scaled
- Large switching/minimal routing
- Distributed routing/switching

Each of these designs follows the three-layer hierarchical structure of a core, distribution, and access layer.

Scaled Switching

Scaled switching is the simplest for configuring in a small campus environment. Each switch is part of a single broadcast domain and therefore uses the same addressing scheme for each connection. Expanding a scaled switching environment is possible through the use of VLANs to segment the network into multiple broadcast domains. Implementing VLANs, however, will require the use of a router for inter-VLAN communications. Figure 7-8 diagrams a scaled switching topology.

Large Switching/Minimal Routing

The large switching/minimal routing topology employs LAN switching at the access layer and either ATM switching or LAN switching at the distribution and core layers. In Figure 7-9, the ATM switching topology requires the following for proper deployment:

- LANE support on all routers and switches.
- Support for ATM UNI 3.x or higher signaling, using point-to-point and point-to-multipoint.

If using LAN switching in the distribution layer, VLAN trunking (IEEE 802.10 or ISL) must be supported in all devices. Additionally, the switches must be running spanning-tree protocol, which inhibits load balancing.

Scaling the large switched/minimal routing topology requires a hierarchical design built on VLANs. The VLAN design should be one that enables the 80/20 rule over the VLANs. This minimizes the use of the routers in the distribution layer.

Figure 7-8
Scaled switching
topology




Figure 7-9
Large switching/
minimal routing
topology

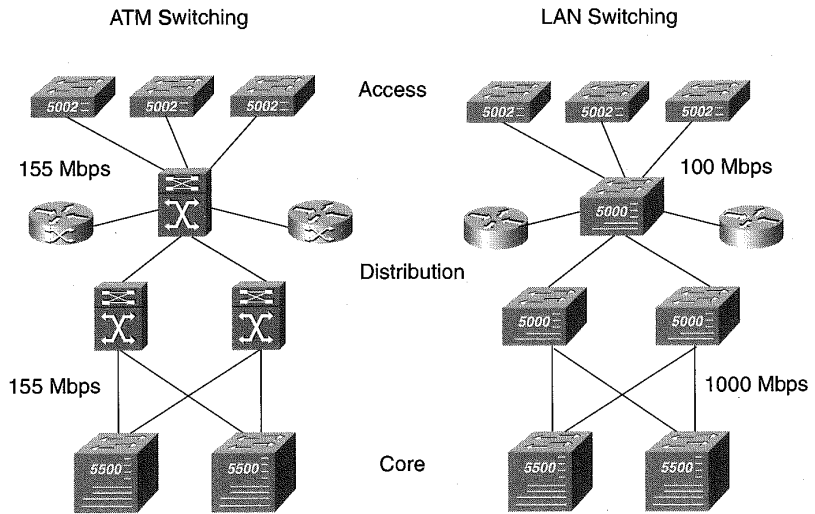



Figure 7-10
Distributed routing
topology



Distributed Routing/Switching


In this topology, the distribution layer is populated with routers only. In either the ATM or LAN switching configuration, shown in Figure 7-10, the routers become the focus of decision making.

This switching topology design follows the classic hierarchical network model. The routers distribute the traffic between the core and the access layers. The use of routers actually serves the new 20/80 rule well. The consolidation of services at a central location is also served well by this design.

CHAPTER

8

SRB/RSRB Network Design



Cisco Systems' support of source-route bridging protocols provides two methodologies specific to using SRB transport. These are pure SRB within a single router and *Remote SRB* (RSRB) for transporting frames using an SRB protocol between routers. An SRB protocol is a direct result of IBM's Token Ring network architecture. Although SRB is designed to bridge Token Ring networks, it is for all intents and purposes a routing protocol, since the Token Ring frame contains the route between the source and destination resource.

Although many Cisco-based router networks still utilize SRB and RSRB methods of transport, the preferred methodology for transporting non-routable protocols over a multiprotocol network is *Data Link Switching* (DLSw). DLSw will be discussed in Chapter 9, "DLSw+ Network Design."

Steps to Effective SRB Design

Cisco's support of SRB network protocol enabled Cisco to become the key networking component for transporting IBM SNA and NetBIOS traffic over a multiprotocol WAN. The experience of transporting these deterministic protocols has led to a "science" for achieving service levels over a multiprotocol network. Effective design of an SRB/RSRB network follows the Cisco hierarchical layer model. Additionally, an effective SRB/RSRB network design also entails the type of protocols supported by SRB/RSRB, link redundancy, controlling explorer packets, and performance factors such as:

- WAN frame sizes
- Managing local acknowledgment
- The selection of the IP routing protocol to minimize convergence time.

Determine SRB Protocols

In a corporate-wide network architecture, there have been essentially two protocols of vital concern, *IBM Systems Network Architecture* (SNA) and *Network Basic Input Output Services* (NetBIOS). IBM SNA carries over 80 percent of worldwide corporate information, and NetBIOS was the first LAN protocol for *personal computer* (PC) communications. SRB transports these protocols between bridges connecting multiple LANs, and RSRB transports these protocols between LAN segments residing in geographically different locations over a WAN.

Determine Parallel Link Requirements

High availability is a typical requirement in any network design. Due to the sensitivity of protocols that use SRB/RSRB for transport, however, parallel and redundant links must be given special attention. The underlying IP routing protocol used over the parallel links and the type of encapsulation methods that are used greatly affect the throughput and timeliness of SRB/RSRB traffic. The capability for the underlying routing protocol to quickly converge after a link failure becomes a critical concern to maintain the SRB sessions during convergence. RSRB, though initially developed for transporting over WANs, enables Cisco routers to transport SNA and NetBIOS over non-SRB networks like Ethernet. Many router network designs include multiple routers connecting to a single Ethernet or Token Ring for high availability and redundancy. In these types of configurations, management and control of explorer frames and broadcasts becomes a challenge to avoid routing loops.

Specify an Appropriate RSRB Encapsulation Technique

Cisco routers employ three encapsulation techniques for transporting SRB protocols between routers. These are:

- Direct
- TCP
- *Fast Sequenced Transport (FST)*

The direct encapsulation technique places the SRB frame into the data field of the media data link frame connecting two routers. Direct encapsulation is used only when there is no intervening router between the source and destination routers.

TCP places the SRB frame into the data field of the TCP segment, which is then transported between the two routers using IP datagrams. TCP encapsulation is used when resequencing the frames becomes a concern prior to delivering the SRB frame to the destination resource.

Finally, the FST encapsulation technique places the SRB frame in the data field of the IP datagram. Since IP does not itself provide for the resequencing of data, it is imperative the FST be used when per-packet load balancing over parallel links is not being used. SRB frames received out of sequence usually end in a disconnect for the SRB session partners.

Cisco Systems' support of source-route bridging protocols provides two methodologies specific to using SRB transport. These are pure SRB within a single router and *Remote SRB* (RSRB) for transporting frames using an SRB protocol between routers. An SRB protocol is a direct result of IBM's Token Ring network architecture. Although SRB is designed to bridge Token Ring networks, it is for all intents and purposes a routing protocol, since the Token Ring frame contains the route between the source and destination resource.

Although many Cisco-based router networks still utilize SRB and RSRB methods of transport, the preferred methodology for transporting non-routable protocols over a multiprotocol network is *Data Link Switching* (DLSw). DLSw will be discussed in Chapter 9, "DLSw+ Network Design."

Steps to Effective SRB Design

Cisco's support of SRB network protocol enabled Cisco to become the key networking component for transporting IBM SNA and NetBIOS traffic over a multiprotocol WAN. The experience of transporting these deterministic protocols has led to a "science" for achieving service levels over a multiprotocol network. Effective design of an SRB/RSRB network follows the Cisco hierarchical layer model. Additionally, an effective SRB/RSRB network design also entails the type of protocols supported by SRB/RSRB, link redundancy, controlling explorer packets, and performance factors such as:

- WAN frame sizes
- Managing local acknowledgment
- The selection of the IP routing protocol to minimize convergence time.

Determine SRB Protocols

In a corporate-wide network architecture, there have been essentially two protocols of vital concern, IBM *Systems Network Architecture* (SNA) and *Network Basic Input Output Services* (NetBIOS). IBM SNA carries over 80 percent of worldwide corporate information, and NetBIOS was the first LAN protocol for *personal computer* (PC) communications. SRB transports these protocols between bridges connecting multiple LANs, and RSRB transports these protocols between LAN segments residing in geographically different locations over a WAN.

Determine Parallel Link Requirements

High availability is a typical requirement in any network design. Due to the sensitivity of protocols that use SRB/RSRB for transport, however, parallel and redundant links must be given special attention. The underlying IP routing protocol used over the parallel links and the type of encapsulation methods that are used greatly affect the throughput and timeliness of SRB/RSRB traffic. The capability for the underlying routing protocol to quickly converge after a link failure becomes a critical concern to maintain the SRB sessions during convergence. RSRB, though initially developed for transporting over WANs, enables Cisco routers to transport SNA and NetBIOS over non-SRB networks like Ethernet. Many router network designs include multiple routers connecting to a single Ethernet or Token Ring for high availability and redundancy. In these types of configurations, management and control of explorer frames and broadcasts becomes a challenge to avoid routing loops.

Specify an Appropriate RSRB Encapsulation Technique

Cisco routers employ three encapsulation techniques for transporting SRB protocols between routers. These are:

- Direct
- TCP
- *Fast Sequenced Transport (FST)*

The direct encapsulation technique places the SRB frame into the data field of the media data link frame connecting two routers. Direct encapsulation is used only when there is no intervening router between the source and destination routers.

TCP places the SRB frame into the data field of the TCP segment, which is then transported between the two routers using IP datagrams. TCP encapsulation is used when resequencing the frames becomes a concern prior to delivering the SRB frame to the destination resource.

Finally, the FST encapsulation technique places the SRB frame in the data field of the IP datagram. Since IP does not itself provide for the resequencing of data, it is imperative the FST be used when per-packet load balancing over parallel links is not being used. SRB frames received out of sequence usually end in a disconnect for the SRB session partners.

Determine WAN Frame Size Requirement

The size of the WAN frame supported on Cisco routers is 4 KB. However, the type of encapsulation technique influences the size of the frame over the WAN. For instance, using TCP encapsulation on the router will fragment large frames to multiples of 1492 bytes. Multiple small frames will be placed into a single 1492-byte frame for transmission over the WAN and will then be disassembled on the receiving side. The importance to understanding this is the amount of buffer space required on the input and output queues within the router.

Determine Local Acknowledgment Requirement

SNA supports three types of acknowledgment frames:

- Receiver-ready
- Receiver-not-ready
- Frame-reject

These frames travel typically between the IBM mainframe communications controller (IBM 3745) and the SNA physical unit (IBM 3174) device. Cisco routers employ a mechanism of responding locally to these SNA acknowledgment frames, thereby minimizing the impact they have on the WAN bandwidth. Local acknowledgment is supported only under RSRB TCP/IP encapsulation.

Select the Appropriate IP Routing Protocol

The underlying routing protocol is a major factor in the capability of an SRB/RSRB network to scale as the demand for the service grows. The key criteria in selecting the underlying IP routing protocol is convergence time, network topology, and routing table maintenance. Basing the design on these criteria will enable a SRB network to scale appropriately.

Explorer Packet Control

SRB network resources use an explorer packet to determine the path to the destination resource. Cisco IOS software in support of SRB networks

employs two methods of managing the explorer traffic. These are *Routing Information Field* (RIF) cache and proxy explorer. The RIF cache maintains a lookup table of MAC addresses attached to the router interfaces that use SRB for transporting traffic. The proxy explorer function places remote MAC addresses in the RIF cache that are for resources communicating with local MAC addresses attached to the router. The proxy explorer function replies back to the requesting station, the cached RIF entry, thereby reducing WAN explorer storms. Proxy explorer in an SNA environment is plausible when the destination MAC address to reach the mainframe is not duplicated on an active target resource.

In a situation where two active destination MAC addresses are found, only the first entry in the RIF cache is used. However, other routers in the network may end up using the second active MAC address. Load balancing to duplicate MAC addresses for mainframe connectivity from a single location is performed by DLSw+ and will be discussed in Chapter 9.

NetBIOS Traffic Management

Traffic management of NetBIOS protocols is a requirement in SRB/RSRB network design due to its inherent broadcast nature. NetBIOS stations are constantly sending out broadcasts to determine the location of partner resources, servers, and to maintain the connection between the requester and the server. Cisco IOS software manages this traffic through specific techniques. NetBIOS name-caching performs a similar function to proxy explorer, but instead of a MAC address and RIF entry, the name cache stores the Mac address, RIF, and NetBIOS name.

Directed named datagram broadcasts are cached in the same manner as the NetBIOS name. NetBIOS stations often send multiple copies of the name_query frame for locating a partner. A Cisco router will throttle the search by sending only a copy of the first frame it received. It will ignore the remaining frames until the query timer expires. The final method for managing the NetBIOS traffic is through dampening. This is similar in use to IP split-horizon. The router will not forward a broadcast to the WAN if the destination of the broadcast is cached to an interface on the router.

Typical SRB Topologies

Three basic SRB topologies exist: hierarchical, distributed, and flat. Each topology has its own characteristics and requirements for implementation.

Hierarchical

As shown in Figure 8-1, the Cisco hierarchical-layered architecture is applied to support an SRB network using the physical connections of the routers. This hierarchy supports the many-to-one architecture found in SNA connectivity to an IBM mainframe. The architecture enables the distribution routers to apply policies in support of class-of-service as well as to protect the access services layer from unwanted frames.

Distributed

In a distributed SRB network with routers, the end user community must communicate with more than one mainframe. In such a scenario, as diagrammed in Figure 8-2, the network forms a dual star configuration, enabling the remote locations to access either of the mainframe sites over the network.

In the distributed topology, end user stations at the remote locations can access either mainframe through the services provided by the core layer of the Cisco hierarchical-layered model.

Flat

The flat topology for an SRB network physically connects the access layer routers with each of the distribution layer routers, shown in Figure 8-3. Although this provides for high availability, it becomes costly to implement and adds to the complexity of route decision processing and troubleshooting. This topology enables any-to-any SRB connectivity.

Figure 8-1
Hierarchical network topology in support of SRB networking on Cisco routers

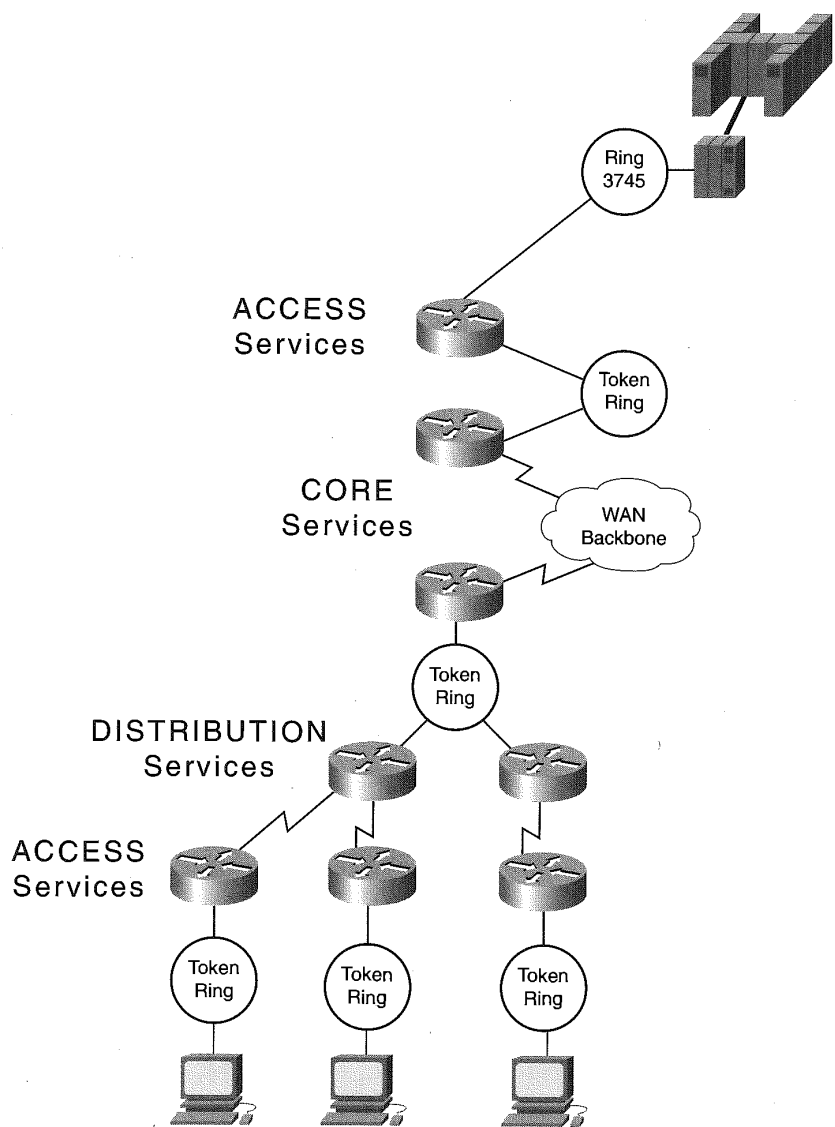


Figure 8-2
Distributed network
topology in support
of SRB networking on
Cisco routers

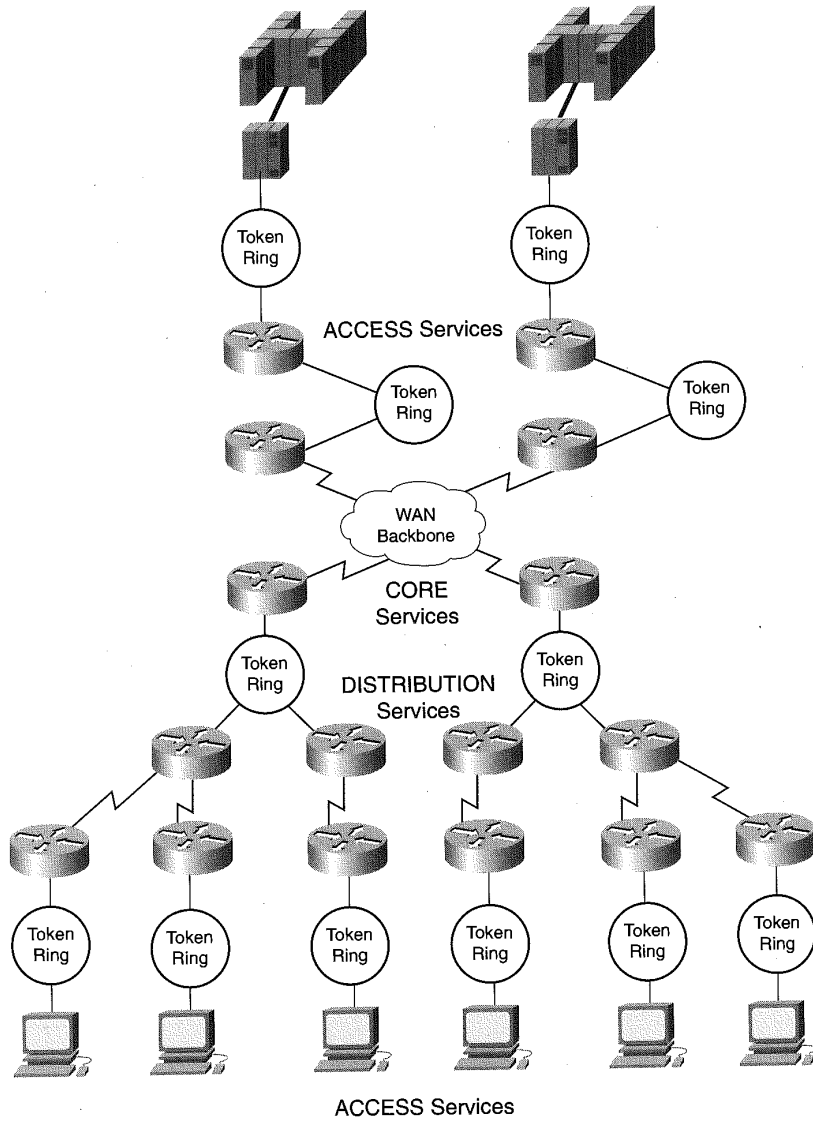
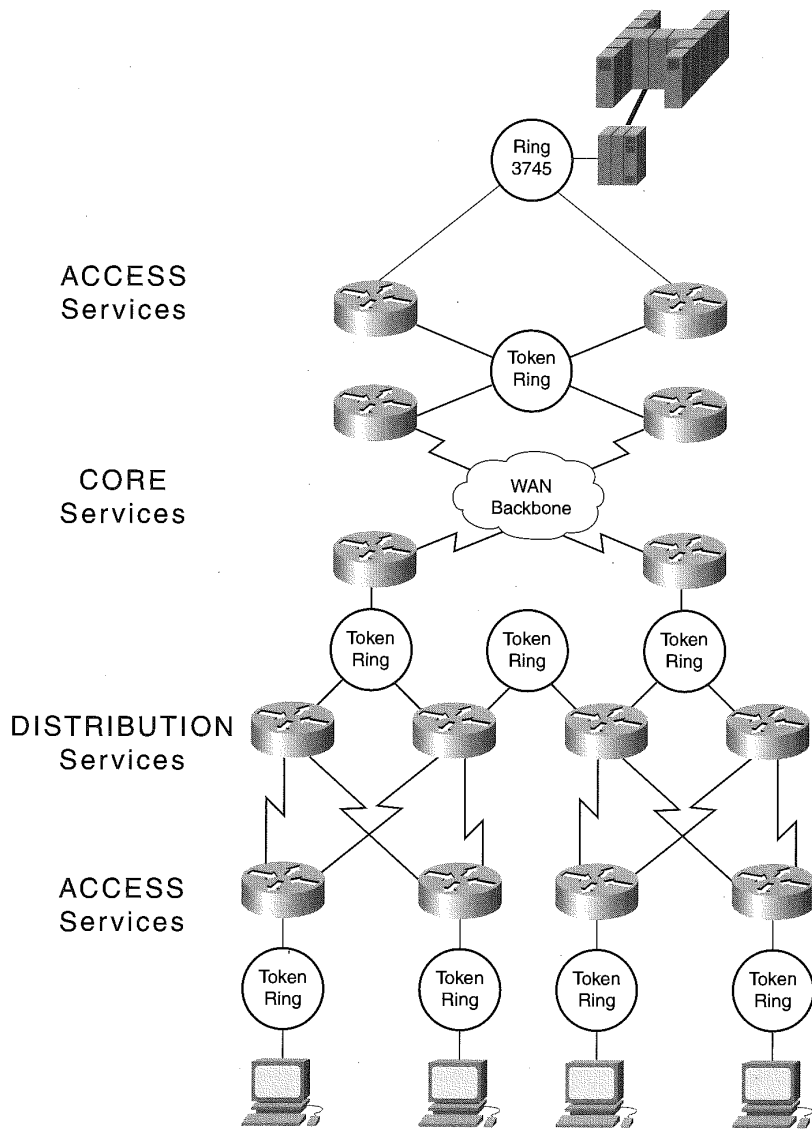



Figure 8-3
Flat network
topology in support
of SRB networking on
Cisco routers



Virtual Ring Concept

Cisco connects remote SRB LAN segments using a virtual Token Ring concept. The virtual ring concept is used as a means of bridging two disjointed SRB networks. Each networking device along the connection between the two remote LAN segments employs and understands the concept of the virtual ring. In Cisco RSRB virtual rings, a ring number is specified to represent the ring in the RIF field of the SRB frame. Figure 8-4 illustrates the virtual ring concept.

Figure 8-4
Cisco virtual
ring concept



In Figure 8-4, two Cisco routers are connected by a WAN. LAN segment A is assigned ring number 1 and LAN segment B is assigned ring number 2. SRB communications is established through a peer relationship between the two routers. The peers in each router are mapped to virtual ring number 3. Connectivity from station A to Station B is established by a RIF field that contains the following information:

```
Ring1-Bridge1-Ring3-Bridge1-Ring2
```

The Ring3 value in the RIF represents the virtual ring defined on each router. The virtual ring is established by the routers connected using RSRB communications. Although this concept enables remote LAN-to-LAN SRB communications, it is still limited to the IEEE 802.5 restriction of 14 segments and 13 hops. The use of virtual rings on the routers can also be applied to the SRB network topology design.

Multiport Bridging

The concept of virtual rings came about from multiport bridging. The purpose of a bridge is to connect two independent LAN segments, while a multiport bridge connects two or more independent LAN segments.

As shown in Figure 8-5, a single router connects four independent LAN segments. Each LAN segment can communicate with the other. This is accomplished by employing a virtual ring on the router that bridges all the LAN ports. The virtual ring becomes the common or backbone ring for connecting the other rings. The RIF field for connecting any ring to any other ring on this router will include the virtual ring defined on the router. This type of configuration is known as local RSRB.

Redundant Star Topology


A redundant star topology is a partial mesh of the virtual ring concept. As shown in Figure 8-6, this design lends itself to providing high availability to the mainframe while enabling growth to the network. The impact of this growth is found in the requirement of having multiple data center routers to service every 15 to 100 remote peers. Using this design does not allow for any-to-any connectivity between the remote locations.

Figure 8-5

Using a virtual ring to connect multiple ring segments on one router

Ring1-Bridge1-Ring5-Bridge2-Ring2
 Ring1-Bridge1-Ring5-Bridge3-Ring3
 Ring1-Bridge1-Ring5-Bridge4-Ring4

Ring2-Bridge2-Ring5-Bridge1-Ring1
 Ring2-Bridge2-Ring5-Bridge3-Ring3
 Ring2-Bridge2-Ring5-Bridge4-Ring4




Ring3-Bridge3-Ring5-Bridge1-Ring1
 Ring3-Bridge3-Ring5-Bridge2-Ring2
 Ring3-Bridge3-Ring5-Bridge4-Ring4

Ring4-Bridge4-Ring5-Bridge1-Ring1
 Ring4-Bridge4-Ring5-Bridge2-Ring2
 Ring4-Bridge4-Ring5-Bridge3-Ring3

Figure 8-6

Partial mesh topology for access to the mainframe using RSRB and virtual rings



Fully Meshed Topology

A fully meshed virtual ring topology is not a scalable design for supporting large SRB networks. A fully meshed design involves the use of one virtual ring used for connecting all the routers in the network. This, while simple in design, proves to be costly for large SRB networks since each router is defined to peer with all the other routers. This type of design will cause excessive CPU utilization and undo the replication of frames to peers that are not the true recipients of the frame. Support for any-to-any communications is best served when employing a hierarchical virtual ring topology.

Hierarchical Virtual Ring Topology

The hierarchical design depicted in Figure 8-7 illustrates the use of virtual rings in a hierarchy. Using this type of topology, each distribution router can handle from 15 to 100 remote peers, while the data center routers can handle the same amount. This type of configuration enables a scalable SRB network to several hundred peers being serviced at the data center for connectivity to the mainframe.

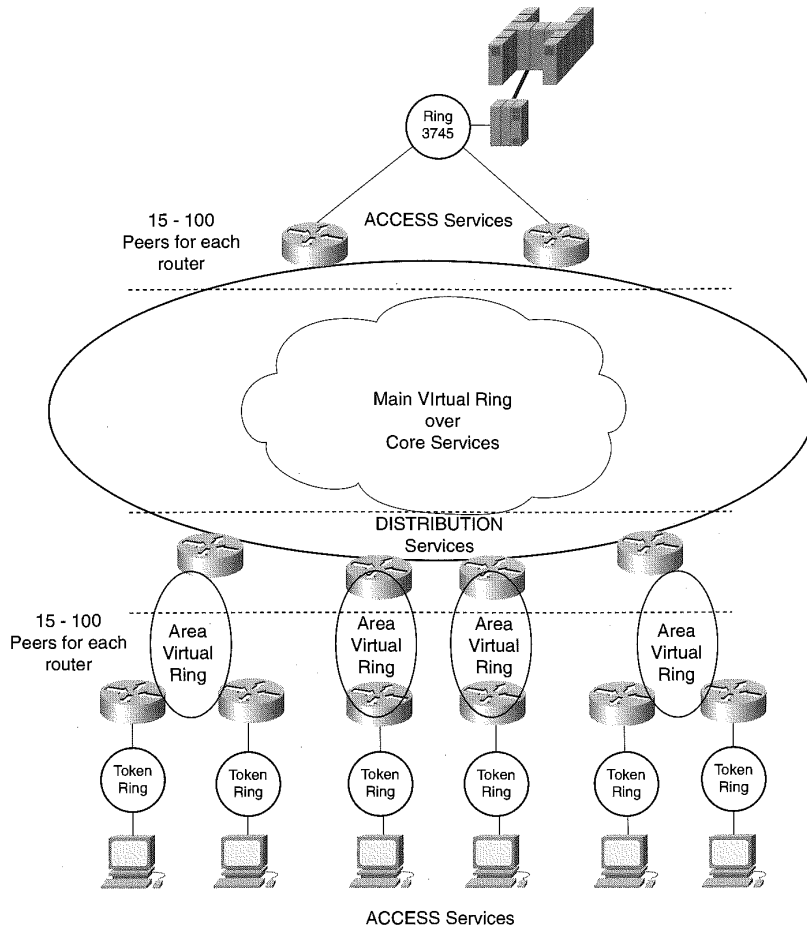
Explorer Packet on Virtual Rings

Although the hierarchical virtual ring topology enables large-scale SRB networks, it does not solve the issue of replicating explorer packets or frames destined for other devices not found on the data center rings. This concern is addressed by applying access filters on the routers to prevent a packet entering a ring if that packet does not originate from the ring that has a FEP MAC address or a known partner MAC address from another remote location. In order to develop this access list, careful analysis of traffic patterns and application usage must be performed prior to creating the filter.

The SRB protocol specifies that a bridge will not forward a frame to a ring that the bridge is connected to if the RIF field of the frame indicates that the frame has already passed through the ring. This rule also applies when using virtual rings. Figure 8-8 illustrates this process.

Figure 8-7


Scalable SRB network topology using hierarchical virtual rings



If multiple virtual rings in the RSRB network design are being defined, an explorer packet will be copied, modified, and then forwarded to the peers associated with the virtual ring on the router. This can cause excessive CPU utilization and could, in fact, depending on the number of peers, cause the SRB network to implode. Cisco IOS allows the network engineer to control this anomaly by creating an explorer queue, whereby end user information frames and bridging frames will be serviced more than explorer frames. The tradeoff of using an explorer queue is the potential of suboptimal paths returned by the explorer packet.

Figure 8-8

Explorer packet
explosion over virtual
rings




Proxy Explorer

A second method of managing the explosion of SRB explorers is the use of the Cisco IOS software feature, proxy explorer. As depicted in Figure 8-9, this feature creates a RIF cache entry for successfully returned explorer packets. The first explorer packet returned ends up being the RIF path used by all subsequent requests.

In Figure 8-9, control unit A is in session with the mainframe MAC address 4000.3745.0001 before control unit B sends an explorer to the same MAC address. The cached RIF entry for the MAC address 4000.3745.0001 is returned to control unit B by the local router. Control unit B then uses the RIF provided by the local router to connect to the mainframe. In this way, the use of proxy explorer minimizes the number of explorer frame explosions for a particular MAC address.

Figure 8-9
Use of proxy explorer
to reduce explorer
explosion



NetBIOS Broadcast Control

NetBIOS is no longer a dominant application protocol for corporate networks, but there may be some instances where NetBIOS or its kin, *Network Basic End User Interface* (NetBEUI), are still in use. These protocols consume large amounts of unproductive broadcast traffic and can therefore swallow up low-bandwidth networks and create broadcast storms over fully meshed high-bandwidth topologies. NetBIOS stations issue broadcasts for the following reasons:

- At startup to verify that a station name is unique within the network (NetBIOS functional address C000)
- To locate a router to a particular server (NetBIOS functional address 0000)
- As a heartbeat function that ensures a viable connection between the client and the server (NetBIOS functional address 0080)

The broadcasts can use either a NetBIOS name or the NetBIOS functional address. Requests sent by NetBIOS stations use a spanning explorer broadcast, while the response is always all-routes explorer packets. The all-routes explorer response means that for every request received through the network from a particular station, a response will be sent. This is compared to a single-route explorer packet response, where only the first request is

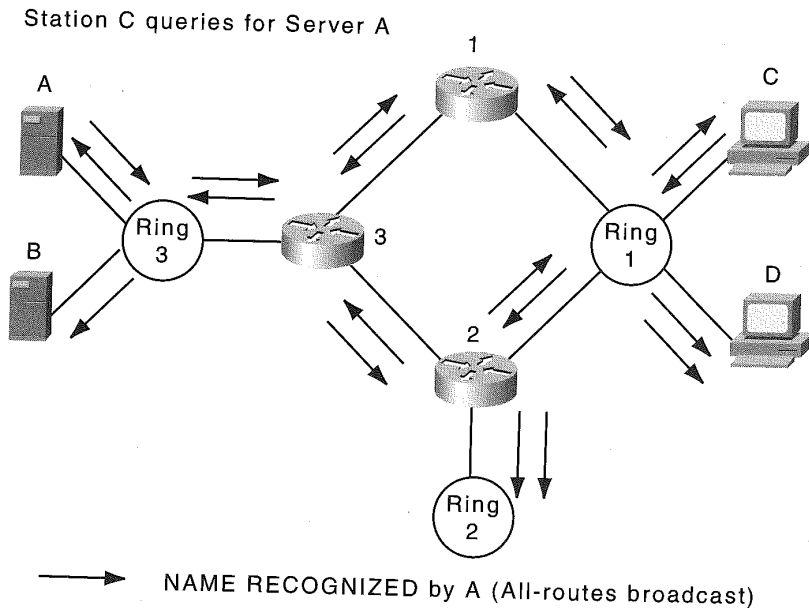
replied while the remaining received explorers are discarded. The subsequent explorers are discarded because the theory of SRB is that the best path between two resources is the receipt of the first explorer.

NetBIOS Name Caching

The caching of NetBIOS names in the routers greatly reduces the NetBIOS explorer process. As shown in Figure 8-10, a NetBIOS NAME-QUERY is passed to each router in the sample network. During this process, the sending station's NetBIOS name, MAC address, and RIF are cached in each router. The response from the named server, in this case server A, enables the routers to cache a NetBIOS name entry for server A in each router.

If station D were to query for server A's routers 1 and 2, having cached the entry for server A from station C's query, it performs a directed query to server A to verify its existence. The reply back from server A through the router network enables the mapping of station D in each router. At this point, only directed queries are performed for server A and any queries for station C and D. This greatly reduces the overhead placed on the network by NetBIOS NAME-QUERY requests.

Figure 8-10
NetBIOS name
caching process
diagram



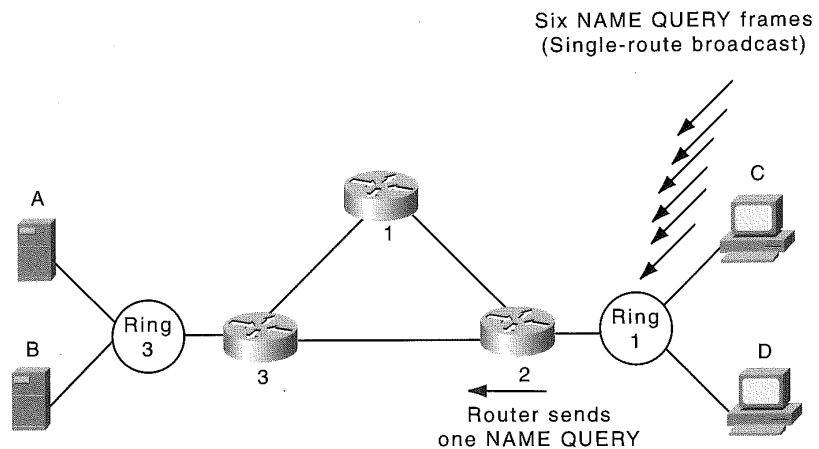
NetBIOS Datagram Broadcast

NetBIOS datagram broadcasts are similar to NetBIOS NAME-QUERY and NAME-RECOGNIZED broadcasts. These broadcasts are also cached in the NetBIOS name cache. The difference between the broadcasts is that the datagram broadcast is usually a one-way broadcast with no corresponding reply.

NetBIOS Broadcast Throttling

The throttling of NetBIOS NAME-QUERY frames eliminates the duplication of frames on the network sent by the source station. Typically, a NetBIOS station will send as many as six copies of the NAME-QUERY frame with a half-second pause between each transmission. The Cisco IOS software enables a throttling of these duplicate frames by discarding subsequent copies of the original frame, as diagrammed in Figure 8-11. The router will only send a subsequent frame from the original station for the same destination name once the dead timer expires. Using this mechanism further reduces NetBIOS protocol overhead on the network.

Figure 8-11
NetBIOS broadcast
throttling



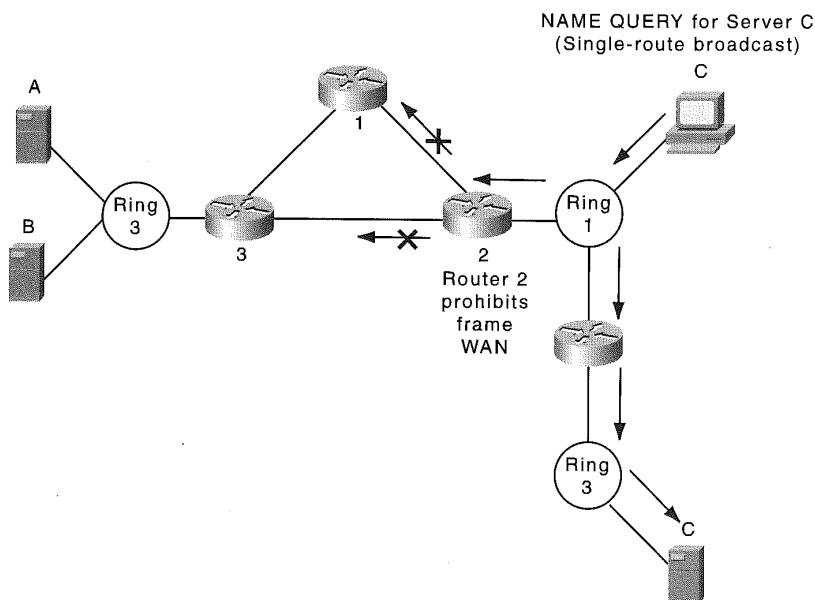
NetBIOS Broadcast Dampening

Dampening of NetBIOS broadcasts is performed by a Cisco router when a broadcast for a resource is found in the name cache table, which contains an entry that points to the receiving interface of the broadcast. Shown in Figure 8-12, station C broadcasts for server C. Router 2 receives the broadcast and inspects the NetBIOS name cache table. The table indicates that the query is for a resource attached to the ring on which the broadcast was received. Router 2 then prevents the broadcast from the WAN.

Remote SRB Encapsulation Techniques

Remote SRB (RSRB) is nothing more than another technique to tunnel non-routable protocols over a router backbone network. Each RSRB tunnel is treated as a virtual Token Ring segment. The virtual ring number used to identify the tunnel is prepended to the received SRB frame as a 16-byte header. The RSRB segments can be associated as a group, thereby having

Figure 8-12
NetBIOS broadcast dampening.




multiple connections on the same virtual Token Ring segment. Each router participating in a RSRB ring group defines the remote routers also participating in the group as peers. The non-routable frames are encapsulated into a frame that can be sent by a router to the WAN or directly to another router over LAN media. RSRB supports three types of encapsulation:

- Direct
- TCP/IP
- FST

Direct

Direct RSRB encapsulation places the frame completely into the data field of the data link frame. As depicted in Figure 8-13, the data link frame can be HDLC for serial-line connectivity, Ethernet, Token Ring, or FDDI encapsulations. Direct RSRB is used only when no intermediate routers are involved with the connection. Direct encapsulation uses the least amount of CPU cycles and frame overhead since it is basically switching the frame from one port to another.

Figure 8-13
Direct RSRB
encapsulation
techniques

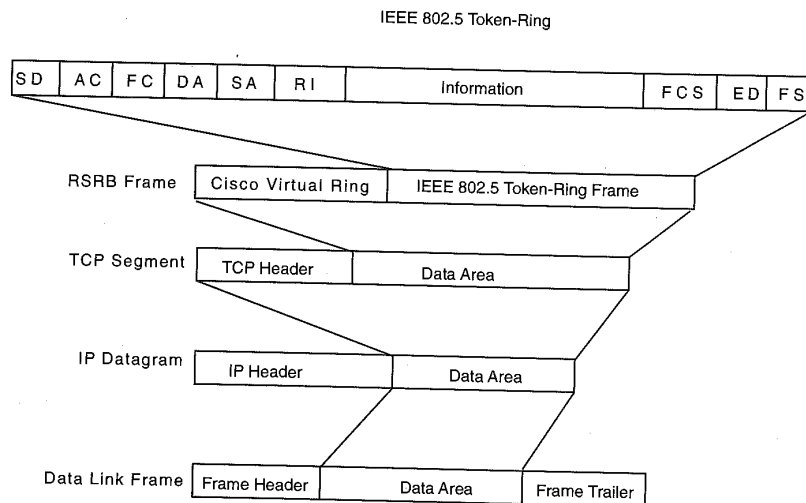


TCP Encapsulation

Using TCP encapsulation for the transport of non-routable protocols over RSRB ensures that the frames are received by the far end device in sequential order. This ensures that the SNA session riding on this RSRB connection does not disconnect its session with VTAM because of out-of-sequence frames, which is an important factor when dealing with the sensitivity of SNA and its basis on point-to-point connections.

TCP encapsulation, shown in Figure 8-14, requires the frame received from the LAN to be placed in the TCP segment data area. RSRB appends the Cisco virtual ring 16-byte header to the IEEE 802.5 Token Ring frame. The TCP segment, including the RSRB frame, is placed into the data area of the IP datagram. The completed IP datagram is then placed in the data area field of the data link frame used for connecting the SNA device to the destination router. This encapsulation technique, while ensuring in-sequence delivery of SNA frames, increases CPU cycles on the router and increases frame overhead due to the TCP and IP header information required.

Figure 8-14
RSRB TCP
encapsulation



Fast Sequenced Transport (FST) Encapsulation

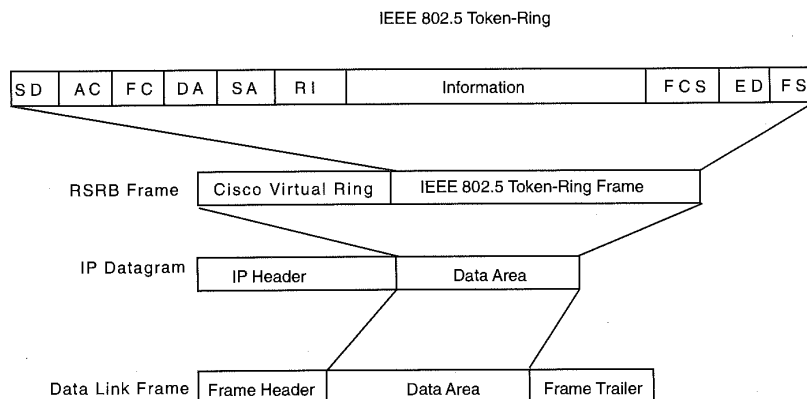
Cisco routers are capable of sending frames encapsulated in IP datagrams, thereby reducing frame overhead and CPU cycles for RSRB encapsulation. This technique is known as *Fast Sequenced Transport (FST)*. CPU cycles are saved because the FST process performs the encapsulation on the interface card. Because IP does not ensure sequenced delivery of frames, FST relies on the LLC2 logic for determining out-of-order frame retransmission. Figure 8-15 illustrates the encapsulation technique.

Care should be taken when employing the FST technique. FST should only be used when you can guarantee that the frames will not be received out of sequence.

Parallel WAN Link Issues

Parallel links between the source and destination pose a design issue for SRB networks. SRB network resources require in-sequence frames. Parallel links of equal cost can allow out-of-sequence frames, which can then force the SRB resource to disconnect the session. Parallel links over a WAN are even more sensitive to this problem because delay factors alone can cause out-of-sequence packets.

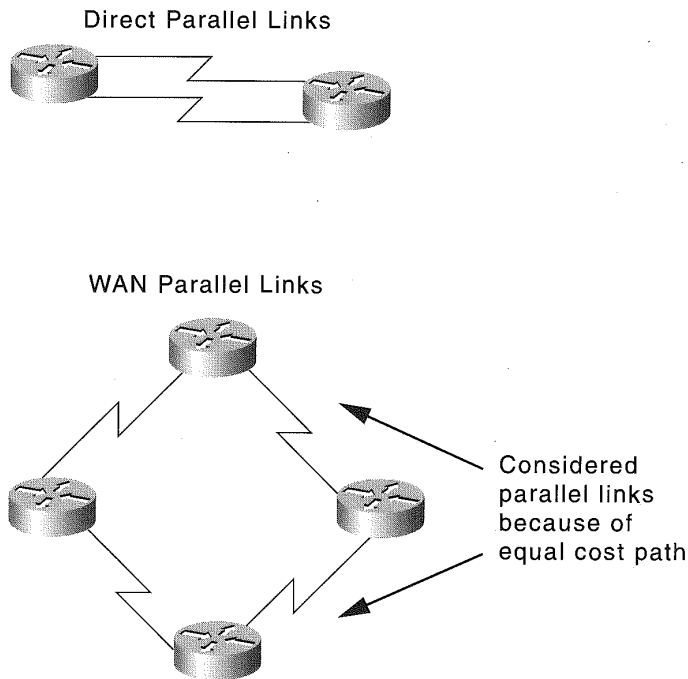
Figure 8-15
RSRB encapsulation
with FST



Parallel links can be two or more dedicated links directly connecting two routers. They can also be physical connections of equal cost that have intermediate network resources. Figure 8-16 diagrams these two configurations.

WAN parallel links can use two different types of switching technologies. Process switching, which takes place in the router processor, performs full evaluation of the routes and per-packet load balancing over the parallel links. Fast switching, using memory on the interface card along with the route processor, performs full route evaluation and per-destination load balancing. Per-packet load balancing, as used by process switching, has the potential of out-of-sequence packets. Per-destination load balancing greatly reduces the potential for out-of-sequence packets.

Figure 8-16
Parallel link
topologies



Process Switching

In process switching, each frame is sent to the route processor CPU, which encapsulates or de-encapsulates the data. In addition, route selection and filtering are performed during process switching. TCP/IP encapsulation requires process switching, which is CPU-intensive; it slows down packet throughput and uses valuable CPU. If you use priority queuing, custom queuing, or filtering, frames become process-switched.

Fast Switching

Fast switching is a method used on Cisco routers when a Cisco bus interface processor is installed. This method enables the passing of a frame that is destined for a port on another bus interface processor directly over the router bus backplane at the interrupt level, without copying the frame to the system memory first. This bypasses the involvement of the route processor CPU and therefore provides better performance.

Fast switching is the default when SRB is enabled. After it is determined that an IP destination can be fast switched, the route is cached and associated with the interface reflected by the route. Direct encapsulation only uses fast switching, while FST encapsulation techniques can utilize fast or process switching. If TCP encapsulation is used, then the TCP connection is process switched.

Effect of IP Routing Protocols

IP-routing protocols widen the parallelism of WAN links through the normal routing algorithms. The routing protocols strive to achieve load balancing and therefore take into account equal-cost paths, equal-cost load balancing, and process switching for each IP destination. The effect of equal-cost paths, as determined by IP-routing protocols, plays heavily on RSRB connections.

The RIP-routing protocol determines parallelism based on equal hop counts to the destination. RIP uses up to four parallel equal-cost paths for load balancing the traffic to an IP destination. IGRP and EIGRP also use up to four equal cost paths in parallel. The metrics are based on bandwidth, delay, reliability, MTU size, and load. OSPF uses its metrics to determine equal cost paths and will use up to four parallel equal-cost paths for load balancing traffic.

IGRP, EIGRP, and OSPF inherently split RSRB traffic across the equal-cost links when the router is using process switching, as is the case with TCP encapsulation. FST encapsulation forces per-destination load balancing and will therefore have all traffic to the remote FST router use a single link of the parallel links.

IGRP and EIGRP support a feature known as variance that enables load balancing over unequal-cost paths. This feature is viable only when there are clearly physical parallel links to the destination point and the higher-cost path enables the positive forward movement of the packet. Although careful analysis of the network topology is required to determine the validity of using the variance feature, it does provide the elimination of convergence due to link outages, since the alternate path is already known and in use.

Local Acknowledgment

SNA local acknowledgment is used only with TCP encapsulation. Proper use of the feature can eliminate WAN overhead and ensure SNA session continuity. Consider the following circumstances for implementing local acknowledgment:

- Expected long network delays to due WAN design, performance, or utilization.
- The network topology includes low-speed links, high utilization on links, or poor quality links.
- Service levels require SNA sessions to remain active during router convergence.
- Minimization of WAN traffic will lead to better performance.
- WAN traffic analysis reveals that 50 percent or more of the WAN traffic is LLC-type data.
- Reduction in WAN costs is accomplished by enabling low-speed links to service remote locations.
- Using TCP/IP encapsulation ensures reliable transport of data.
- Unreliable WAN links create frequent SNA session losses.
- The modification to retry timers on remote devices becomes an administrative issue.
- Bandwidth constraints require the elimination of acknowledgment traffic.

Design Recommendations

When the network design clearly uses parallel WAN links, the following recommendations should be considered:

- Design the network with routers servicing the SRB stations along with minimal WAN connections to reduce CPU bandwidth requirements.
- FST encapsulation should be used as much as possible.
- TCP/IP encapsulation should be used if local acknowledgment or prioritization is required.
- Maximize fast switching.
- Use TCP/IP encapsulation with IGRP variance in meshed topologies for low-speed links where local acknowledgment and or prioritization is a requirement when the topology can take advantage of these features.

When link speeds are primarily greater than 64 kbps and local acknowledgment is a requirement, follow these recommendations:

- Use TCP/IP encapsulation only on those links that have a history of session loss (local acknowledgment) and when high-speed links are servicing the location.
- Use FST encapsulation on high-speed links when local acknowledgment or prioritization is not a requirement.

IP Routing Protocols and SRB

The type of routing protocol employed within the network and its capability to converge rapidly greatly affects SRB connections. SNA sessions are time-sensitive and an SNA device will cancel its session if connectivity has not been confirmed in a certain amount of time.

Typically, without the use of local acknowledgment, an SNA connector will fail anywhere from 13 to 42 seconds once a link failure has occurred on the route being used for transporting the SNA session data. The time to converge must fall within a period that satisfies the shortest inactivity timer for an SNA device using the network. This will ensure SNA session survivability during link outages.

Link Failure Effects

The detection of link failures is the initial determining factor in reducing convergence time. Cisco IOS provides parameters that allow a network engineer to specify values that can reduce the detection time and thereby reduce the overall convergence time. Although in general the detection of link carrier-loss means a failed link, poorly performing links can also indicate a failed link. For example, a connection at a remote location may have gone inoperable, but at the local link connection, the carrier indicator is active. Routers use keepalive and hello packets to determine the operational stability of a link.

Serial link connections are the most unreliable of links. A serial interface has a default keepalive packet sent over the link every 10 seconds. The router determines that a link is inactive if a response from the keepalive has not been received after three keepalives have been missed. This means that failed link detection could take up to 30 seconds. Modifying the keepalive timer to three-second intervals realizes a failed link detection within seven to nine seconds.

Token Ring link connections inherently determine failed link circumstances immediately since the Token Ring network is disrupted. Keepalives can be used to provide a mechanism that not only determines a failed connection on the Token Ring network, but if the ring is under a heavy load.

FDDI link connections detect failure immediately. FDDI is a highly reliable media and, as such, the router will inactivate an FDDI interface under exceedingly high-traffic volumes. The most common cause of FDDI failure is turning off a device that is attached to the FDDI ring, causing the ring to wrap in a dual-attached FDDI environment.

Ethernet itself is fairly reliable, but its basis on CSMA/CD leaves it with a poor link failure detection mechanism. Keepalive packets on Ethernet links therefore are the only means for accurately recognizing a link failure. As with serial-line interfaces, the keepalive defaults to 10 seconds and must fail in getting a response in three tries before the router disables the Ethernet interface. Failure detection on Ethernet can take anywhere from 21 to 30 seconds using the default 10-second keepalive timer.

IGRP Network Design

IGRP is a distance-vector-routing protocol and approximates the network topology. As such, it uses the hold down concept to determine whether alternate paths are alternate paths. The hold down timer delays the convergence

of using a new route to the destination for the period specified as hold down. The default values for IGRP timers are inappropriate for SRB networks, and the IGRP update timer defaults to 90 seconds. This timer should be changed to 20 seconds, which reduces the IGRP convergence time to 10 seconds. If hold down is enabled, then at worst case, the total time for convergence is three update periods of 20 seconds when the hold down default of 60 seconds is used. In such a case, local acknowledgment will keep the SNA session up during IGRP convergence.

Enhanced IGRP Network Design

EIGRP has its roots in IGRP and distance-vector routing. Therefore, it uses the same metrics as IGRP, but its implementation ensures a loop-free network-routing table. As such, EIGRP updates only the routers that are affected by the link outage, versus all the routers converging to build new tables.

OSPF Network Design

OSPF also uses the loss of carriers or keepalives on links for failed connections. Additionally, OSPF will use the failure of transmitting or receiving hello packets within a specified period of time. This timer on OSPF is known as the dead timer. Once a link is determined to be down due to the dead timer, exploring an OSPF router will produce an OSPF area-wide broadcast, causing all other OSPF routers to recompute their routing tables. The OSPF dead timer default value is 40 seconds and is too long to support SNA and NetBIOS connections. The standard is to set the OSPF dead timer equal to three times that of the OSPF hello packet time. In support of SRB connections over an OSPF router backbone, it is recommended that the dead timer be set to 18 seconds and the hello timer be set to six seconds for each interface.

Queuing and Prioritization

SNA is a highly deterministic protocol. It has built-in functions for delivering a *quality of service* (QoS) for SNA sessions. Cisco IOS features queuing and prioritization to enable the routers to carry the SNA *class-of-service*

(COS) requirements through the router network. These features are not all specific to SNA and NetBIOS and can be used with other networking protocols.

Priority Queuing

Priority queuing is used to prioritize packets. Enabling priority queuing on a router interface creates up to four output queues for the interface. The queues are separated into low, medium, normal, and high. The router first services packets out of the high priority queue before moving on to normal, then medium, and then low.

In networks where low-speed links are in use, this queuing mechanism can cause lower queued packets never to be serviced, which can result in connection timeouts. This is especially of concern with SNA packets. In such a network, if priority queuing is in use, ensure that SNA packets have the highest queuing priority. Priority queuing has a detrimental effect on router CPU utilization since priority queuing packets must be process-switched.

Custom Queuing

Custom queuing allows a percentage of bandwidth to be given to a particular protocol. Enabling custom queuing results in a total of 16 output queues numbered 0 to 15 for the interface. The router cycles through queues 1 to 15 sequentially. The router does not move on to the next queue until the current queue is depleted or the specified number of bytes to be serviced on the queue prior to moving to the next queue has been sent.

The queue numbered 0 is the system queue. High-priority packets, such as keepalives, are kept in this queue and serviced prior to any other queue. In networks where SNA is the important protocol, using the custom queue mechanism with a high byte count for RSRB protocol will ensure a high level of service for the time-sensitive SNA traffic.

SAP Prioritization

Using SAP prioritization, RSRB protocols can be given precedence over other protocols, based on the *destination service access point* (DSAP) and the *source service access point* (SSAP) addresses. SAP prioritization is used

in conjunction with either priority or custom queuing. Using priority queuing, a precedence can be set for the RSRB SAP address defined. With custom queuing, both precedence and bandwidth can be realized for the RSRB SAP addresses. As an example, SNA traffic using a SAP of 04 can be given preference to SNA traffic using a SAP 08 address.

LU Address Prioritization

Further prioritization of traffic is possible using a specific SNA LU address prioritization technique. Using this mechanism, an SNA physical unit and associated SNA logical units can be given a specific priority. This is important for connections that support both interactive and print traffic. Using the local LU prioritization feature allows interactive traffic to take precedence over print traffic.



SAP Filters for WAN Links

SAP filters can be applied to WAN links to prevent NetBIOS and other LAN protocol broadcasts from traversing the RSRB WAN. Using SAP addresses, SAP filters can be used regardless of the media or encapsulation type. In addition, SAP filters enable the filtering of NetBIOS packets by the NetBIOS name.

CHAPTER

9

DLsw+ Network Design



Data link switching (DLSw) was defined by IBM on its introduction of the IBM 6611 router as a means of transporting SNA and NetBIOS traffic across TCP/IP-based backbone networks. The definition was submitted by IBM to become a standard and was accepted as RFC 1434. This first incarnation of DLSw included TCP encapsulation with local acknowledgment for transporting data. RFC 1434 also overcame the hop limits of source route bridging by massaging the RIF field of the packet to emulate one hop to the destination.

Since RFC 1434, a second RFC that obsoletes RFC 1434 was introduced and accepted. RFC 1795 enhances DLSw by including the prioritization of packets and terminating the RIF field at the virtual ring. Termination of the RIF at each router's virtual ring allows the RIF to extend to its fullest implementation of seven hops or 13 hops (IEEE 802.5) on the downstream side of the router. Finally, RFC 1795 enabled the DLSw partners to exchange information for divulging the capabilities of each DLSw partner, alleviating configuration errors.

A third iteration of DLSw has been introduced by Cisco Systems into the RFC standards that enables DLSw networks to scale into large networks, provide enhanced availability, reduce configuration requirements, and increase performance through custom queuing and load balancing. This iteration is commonly called DLSw+. Table 9-1 lists the functional comparisons of RSRB and DLSw versions.

DLSw Standard

The DLSw standard implements a *switch-to-switch protocol (SSP)* that is used between routers. The routers are termed data-link switches in SSP and perform the following functions

- Establish DLSw peer connections using TCP
- Perform a locate function for network resources
- Forward data between source and destination
- Manage flow control between the peer connections
- Provide error recovery on behalf of the source and destination partners
- Automatically perform local acknowledgment
- Multiplex traffic of multiple data link controls over TCP sessions
- Provide reliable transport of SNA and NetBIOS data over an IP backbone

Table 9-1

DLSw Functional Comparison to RSRB

| | Feature | RSRB | RFC1434 | RFC1795 | DLSw+ |
|--------------------------|----------------------------|-----------|----------|------------|------------|
| Transport | TCP w/loc-ack | ■ | ■ | ■ | ■ |
| | FST | | ■ | | ■ |
| | Direct | | ■ | | ■ |
| Performance | Custom queuing | | ■ | | ■ |
| | Prioritization | | ■ | ■ | ■ |
| | Load balancing | | ■ | | ■ |
| Administrative | Capabilities exchange | | ■ | ■ | ■ |
| | Peer costs | | | | ■ |
| | Dynamic peers | | | | ■ |
| | Ring lists | | | | ■ |
| Scalability | Caching | | ■ | | ■ |
| | Hop reduction | Pass-thru | ■ | Terminates | Terminates |
| | Peer/border groups | | | | ■ |
| | On-demand peers | | | | ■ |
| | Broadcast firewall | | | | ■ |
| Availability Optional | Local termination | Optional | Required | Required | Optional |
| | Backup peers | | | | ■ |
| | Directed verify | | | | ■ |
| Management | Std. MIB/per circuit mgmt. | | | ■ | ■ |

The DLSw standard requires DLSw-to-DLSw connections to be established prior to the transmission of end user data. This connection is referred to as a peer connection and each router is considered a peer or partner. Once a peer connection is made, the peers exchange their capabilities in support of the peer connection. Finally, the DLSw peers establish the complete circuit from source to destination network resources.

Peer Connections

Under DLSw RFC 1795, two TCP connections for DLSw are established between routers used for the DLSw connection. One connection is used for sending and the second is used for receiving. RFC 1795 allows for one the connections to be dropped. In this case, the single connection performs both send and receive operations. Typically, this is up to each vendor's implementation of the standard. Additional TCP connections can be created in support of establishing priority levels for traffic between the peers.

Cisco routers implementing the RFC 1795 standard drop one of the TCP connections after establishing the peer connection. If the capabilities exchange indicates the partner is not a Cisco router, however, both TCP connections will remain active.

Exchange of Capabilities

A capabilities exchange occurs after the successful establishment of two TCP connections. This exchange typically involves the following information:

- The DLSw version number (indicating which standard is being supported by the peer)
- The initial receive window size to establish pacing windows for flow control
- If NetBIOS is being supported
- A list of supported *link service access points* (LSAPs)
- The number of supported TCP connections (two connections are required or multiple in support of prioritization)
- The MAC address or NetBIOS name lists defined as potential session level partners to avoid broadcasts
- Limit the type of search frames that can be sent to the partner.

Once this information is transferred between the two DLSw partners, they can support a complete circuit between SNA or NetBIOS network resources over the IP-based network.

Circuit Establishment

The establishment of a circuit involves several steps. First, the DLSw partners required for the end-to-end session must have a peer connection. Second, an end system, such as an SNA physical unit or a NetBIOS workstation, must send an appropriate request to locate the partner end system. For SNA resources, this is either an SNA TEST or XID explorer frame that includes the MAC address of the partner end system.

NetBIOS stations sends the NAME-QUERY frame that includes the NetBIOS name of the destination end system. A DLSw router receives these frames and sends a DLSw CANUREACH frame to each of its active DLSw peer connections searching for the SNA end system MAC address. The DLSw router sends a NetBIOS NAME-QUERY frame containing the target NetBIOS name for a NetBIOS connection. If one of the DLSw peers locates the destination end system, it sends an ICANREACH frame back to the originating DLSw peer for SNA searches and a NetBIOS NAME-RECOGNIZED frame for NetBIOS searches.

At this point, the originating DLSw router issues a circuit setup request to complete the end-to-end circuit for this SNA or NetBIOS session. As the information gathered by the circuit establishment process is learned, each DLSw router caches the MAC addresses and NetBIOS names to eliminate subsequent future explorer frames sent by end systems.

The established circuit is now comprised of the connections between the end systems, their attached router, and the TCP peer connections between the DLSw partners. Each circuit has a unique ID that enables a TCP peer connection to support multiple circuits. The circuit ID is defined using the source and destination MAC addresses, the source and destination LSAPs, and a data link control port ID assigned by DLSw. Once the circuit is established, data can flow between the end systems.

Flow Control

DLSw uses adaptive pacing mechanisms for flow control. It specifically uses two independent, unidirectional circuit flow control mechanisms on a per-circuit basis. These mechanisms use a dynamic window based on buffer

availability, TCP transmit queue depth, and end station flow-control mechanisms. The DLSw standard allows the windows to be incremented, decremented, halved, or reset to zero. The receiver indicates the flow control through use of the following methods:

- *Repeat*: Increments the window by the size of the current window
- *Increment*: Increases the window size by one and modifies the total window size by the result
- *Decrement*: Decreases the window size by one but increases the total window size by one
- *Reset*: The window size is set to zero to stop all transmissions in one direction until a flow-control indicator increments the window
- *Half*: Reduces the current window size by half but increases the overall queue by the halved number

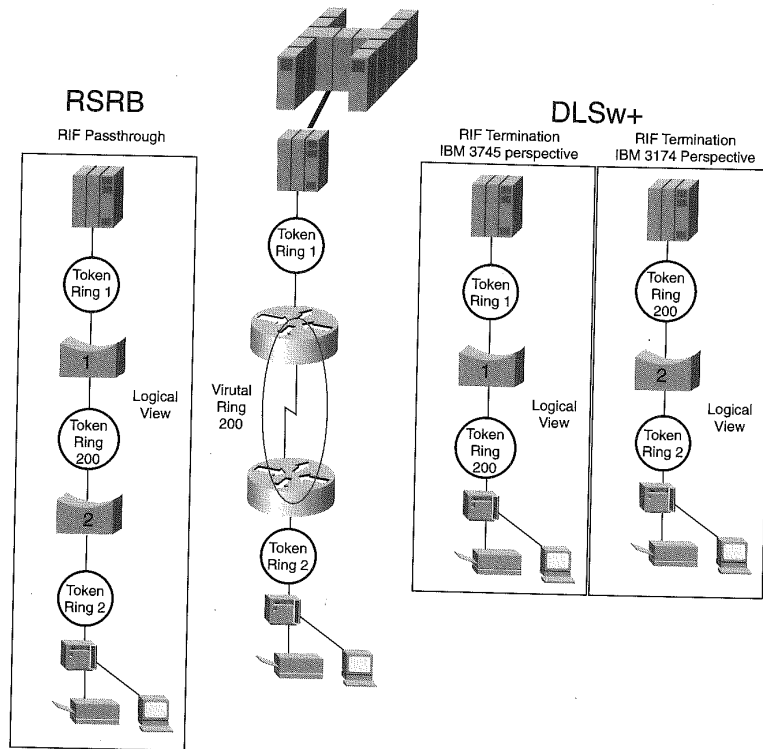
Only the reset indicator must flow in an independent flow-control message. All other flow-control indicators and acknowledgments can be included in the information frames containing end system data or they can flow in their own messages

Cisco DLSw+

The Cisco Systems implementation of DLSw is called DLSw+. The DLSw+ feature builds on the strengths of RSRB while using the scalability advantages of DLSw. At the outset of DLSw+, it was imperative that it be compatible with RSRB, and IOS Release 11.0 supports both RSRB and DLSw+ concurrently. A main reason for this compatibility is to allow large RSRB networks to methodically migrate to DLSw+. The peers used in the connections must be configured as either RSRB or DLSw+ peers. In order for a DLSw+ router to be compatible with RSRB, the peer is defined as an RSRB peer. This allows the DLSw+ router to use the RSRB peers and RSRB to function.

In RSRB networks, the RIF is passed through the WAN, providing full RIF visibility at any point in the connection. Although this is an advantage for troubleshooting using tools like Network Associates SNIFFER, it limits the hop count of the SRB network to 7 or 13, depending on which standard is being implemented for SRB. Using a terminated configuration, like that used with DLSw+, shown in Figure 9-1, the RIF ends at the virtual ring number attaching the routers. This allows the network to scale better than

Figure 9-1
RSRB RIF passthrough
compared to DLSw+
terminating the RIF



RSRB, providing the full seven or 13 hop count limit on the end user side of the routers involved. The downfall to local termination is the inability to view the entire RIF used for the complete connection using tools like Network Associates SNIFFER.

The two encapsulation techniques for reliably transporting SNA data have many features in common:

- TCP, *fast sequence transport* (FST), and direct encapsulation
- FST and direct for Token Ring to Token Ring only
- Access list filtering on LAN and WAN
- WAN prioritization and custom queuing

However, RSRB has some features not found in DLSw+:

- FST and direct encapsulation for SDLLC
- DLC passthrough for TCP encapsulations
- End-to-end RIF passthrough

Many Cisco router-based networks transporting SNA and NetBIOS frames were first established using Cisco *Remote Source-Route Bridging* (RSRB) connectivity. Cisco DLSw+ has incorporated many of the enhanced features of RSRB into the base DLSw+ configuration. Migrating from an RSRB network to a DLSw+ requires the removal of the following functions from your IOS configuration to avoid conflict with the inherent functions of DLSw+:

- Proxy Explorer
- NetBIOS name caching
- SDLC-LLC2 conversion (SDLLC)
- Source-Route/Translational Bridging (SR/TLB)

DLSw+-enabled routers maintain a cache of multiple paths that define the reachability of destination MAC addresses and NetBIOS names. If there are multiple paths to a resource, the path with the least cost is preferred and the remaining paths are labeled capable. Should the preferred path become unavailable, the next path in the reachability path table for the resource becomes the new preferred path. The path entries are stored for remote peers as capable and for local resources as ports. The status capable indicates that the peer associated with the entry at one time was capable of reaching the destination resource. Entries using ports identify that the destination resource is directly attached to the router through one of the communication ports. The multiple paths available for establishing connections are categorized as either fault-tolerant or load-balancing.


Peer Group Concept

DLSw requires all DLSw routers to peer with each other. This requirement comes from the need to facilitate any-to-any communications between end systems. A fully meshed DLSw network, as shown in Figure 9-2, wastes costly bandwidth on explorer traffic between the DLSw peers.

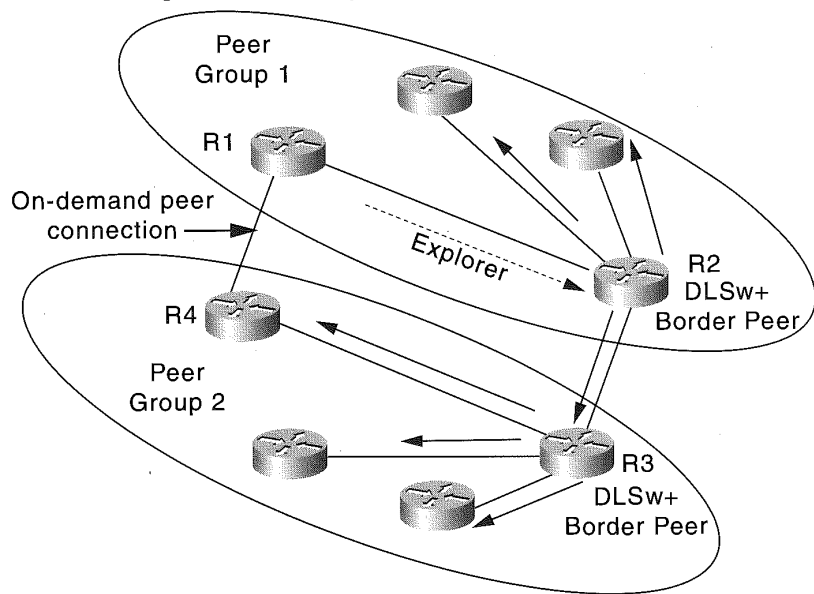
Cisco DLSw+ allows the network designer to logically group the DLSw+ routers into peer groups. Such peer groups address the explorer replication found in fully meshed DLSw+ networks by defining at least one DLSw+ router as a border peer. Each DLSw+ router within the peer group has an

Figure 9-2
The peer group
concept

Logical Full Mesh DLSw+ Peer Network



Logical Peer Groups within DLSw+ Peer Network



active peer connection only with the border peer. The border peer therefore becomes the master MAC address and NetBIOS NAME cache for the peer group. If a DLSw+ router (R1) receives a TEST or NetBIOS NAME-QUERY frame for a resource not found in its own cache, it sends only one explorer to the border peer. The border peer (R2) then checks its cache to locate the resource. If the resource is not in the border peer cache, it forwards an explorer on all its peer connections on behalf of the original DLSw+ router. One of these peer connections is another border peer (R3) that services another peer group. Border peer R3 checks its cache and if the requested resource is not found, it forwards the explorer to the peer group members with which it has peer connections.

In this example, router R4 has found the requested resource and will establish a direct peer connection with router R1 in peer group1 to establish the end-to-end circuit. This connection is termed on-demand peering. The on-demand peer connection remains active while data is being transmitted across the circuit. Once data ceases to be transmitted the peer connection is torn down. This example shows how peer group topology enables the capability of DLSw+ to scale large any-to-any communications by minimizing explorer traffic without persistent TCP peer connections between all DLSw+ routers.

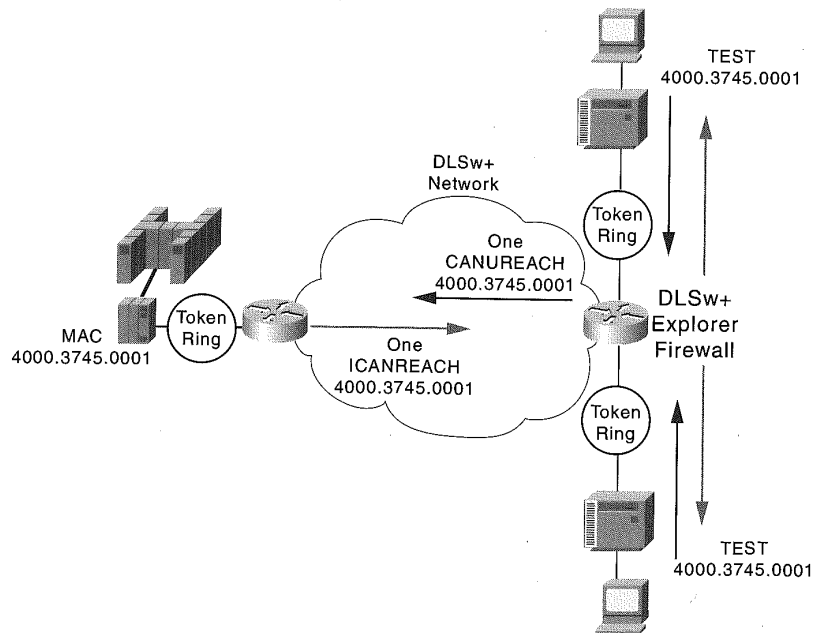
Explorer Firewalls

DLSw+ routers can further protect large networks from explorer storms through the use of the explorer firewall feature. This feature, as depicted in Figure 9-3, forwards only one explorer into the DLSw network, caching other explorer requests for the same destination. Using this feature, normal network occurrences, such as the beginning of the day sessions with IBM SNA mainframes, protects non-DLSw traffic from being overrun by DLSw explorer frames. Once a reply from the explorer is received, the originating DLSw router immediately notifies the end systems that have had their explorers cached of the existence of the destination resource. The DLSw router then attempts to establish a circuit for all the requesting end systems with the found destination partner.

Fault Tolerance

DLSw+ defaults the path as a fault-tolerant path. Using fault-tolerant mode, the originating DLSw+ router sends a response back to the requesting end station from its cache if the entry has not timed-out. If the cached

Figure 9-3
DLSw+ explorer
firewall



entry has timed-out, the originating DLSw+ router sends a CANUREACH request to each of its DLSw+ peers. Any of the peers not responding are removed from the cache.

The destination DLSw+ router responds to a CANUREACH request immediately if the entry in the cache has not timed-out. If the cache entry in the destination DLSw+ router has timed-out, it forwards a single route broadcast test frame over all the ports known to the cache. If any of the ports reply to the test frame, the destination DLSw+ router replies back to the originating DLSw+ router with an ICANREACH response. The originating DLSw+ router then updates its cache and establishes the circuit for the requesting end station.

Load Balancing

Duplicate paths to destination resources allow DLSw+ to be configured for load balancing. This is particularly advantageous for connectivity to IBM hosts that use duplicate MAC addresses for connecting to the same IBM host. The duplicate MAC address is assigned to a LAN interface of an IBM FEP or Cisco CIP internal LAN. The LAN adapters represented by

the MAC address must, however, be connected to different Token Ring segments for them both to be active at the same time.

DLSw+ recognizes duplicate MAC addresses and load-balances the circuits from the end stations requesting connectivity to the mainframe using a round-robin schema based on the cached entry list in the originating DLSw+ router. Figure 9-4 illustrates the load-balancing feature of DLSw+.

The workstation connecting through router R2 in Figure 9-4 establishes its SNA connection through FEP A. Next, the workstation connecting through router R3 sends a test frame looking for the FEP attaching to the mainframe. R1 receives the test frame and inspects the circuit cache for a duplicate MAC address. In this case, a duplicate is found and the R1 router then determines the number of active circuits to each FEP. In this scenario, FEP A has an active circuit and FEP B does not. Using a load-balancing router R1 establishes the circuit to FEP B. The workstation connecting through router R3 is then assigned the path using FEP B by router R1.

Cisco DLSw+ Transport

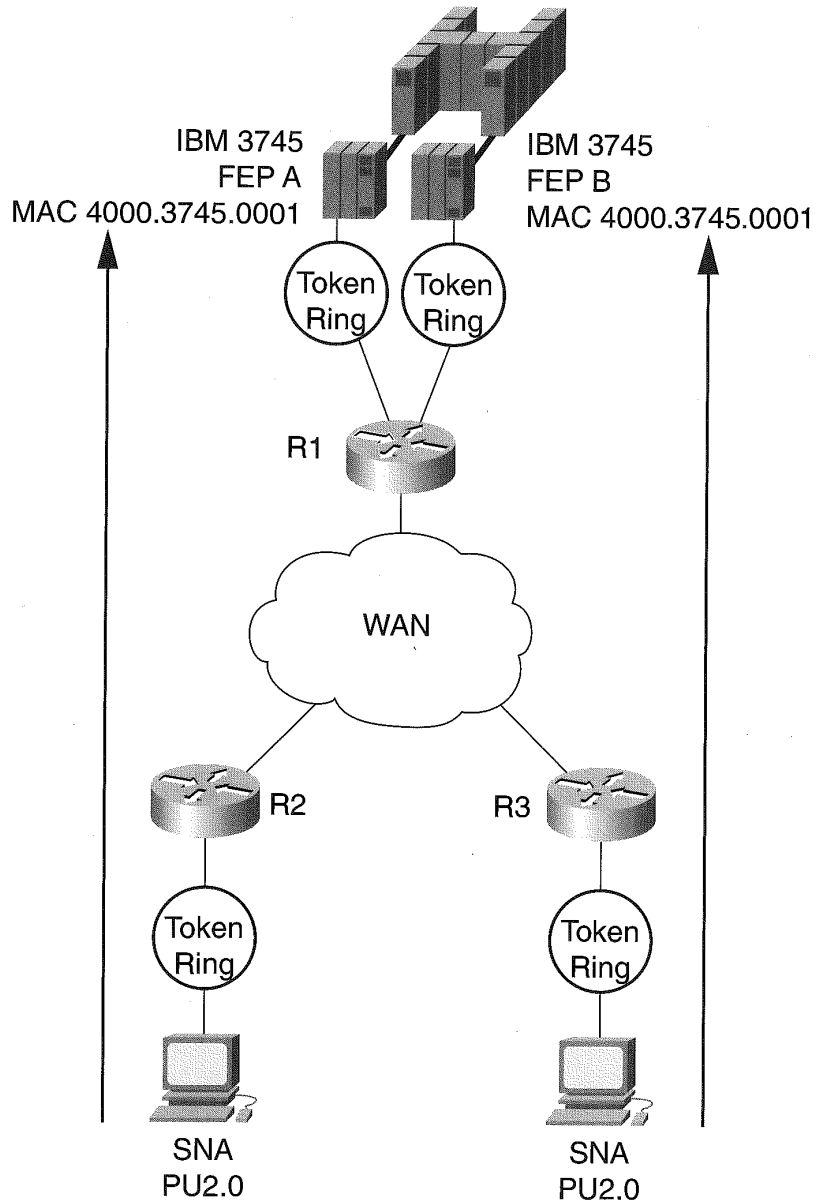
The DLSw standard allows for only TCP connections to carry DLSw. The Cisco DLSw+ implementation of DLSw allows the same type of encapsulation techniques used with RSRB under DLSw+. These are TCP, FST, and Direct encapsulation. In addition, DLSw+ offers an encapsulation technique specific to use on frame relay networks called DLSw+ Lite.

TCP Encapsulation

TCP encapsulation provides reliable delivery of frames with local acknowledgment. Using TCP offers the highest availability of the various encapsulation techniques. TCP is process-switched and therefore uses more CPU cycles than the other techniques. TCP encapsulation adds 20 bytes for TCP header information and 20 bytes of IP header information along with the DLSw header of 16 bytes. TCP header or payload compression can be used to reduce the overhead, but this has to be weighed against the processing time of compressing and decompressing as well as its impact on router CPU and overall response times.

TCP encapsulation virtually eliminates polling overhead through the WAN, due to it using local acknowledgment and local polling. An added benefit of TCP encapsulation is the capability to load balance the DLSw traffic over multiple links, since TCP resequences frames prior to forwarding them to the

Figure 9-4
Load-balancing SNA
sessions using
duplicate MAC
addresses on IBM
3745 FEPs



destination. Since TCP guarantees delivery of frames, it non-disruptively reroutes traffic due to network disruptions.

Another attractive feature of TCP encapsulation with DLSw+ is the capability to prioritize traffic by associating the circuits with different TCP ports. The traffic can also be prioritized by LLC2 SAP addressing, SNA logical unit device addresses, or a MAC/SAP pairing.

Fast Sequenced Transport (FST)

Fast Sequenced Transport (FST) is primarily reserved for use on high-speed links greater than 256Kbps. Although FST with DLSw+ ensures sequenced delivery, it does not guarantee delivery, as does TCP encapsulation. However, the reduced encapsulation mechanism of placing the frame directly in an IP datagram, saving the TCP header encapsulation overhead of 20 bytes, means DLSw+ can process more packets per second than with TCP encapsulation.

The downside to FST is that it does not perform local acknowledgment. Hence, all polling flows and keepalive messages flow over the IP backbone between the end systems. This overhead must be weighed against the process savings. The prioritization schemes allowed by TCP are not available with FST. Therefore, queuing schemes are not as granular.

Direct

Using direct encapsulation for DLSw+ has the same requirement as RSRB. There can be no intermediate routers involved in the connection. The connection must therefore be distinguished as a point-to-point connection. Because of this requirement, direct encapsulation is supported only for HDLC and Frame Relay serial link connections. Although direct encapsulation includes only the 16-byte DLSw header within the HDLC or frame relay frame, it does not provide reliable delivery or local acknowledgment. Hence, disruptive routing is a given and all polling and keepalive messages flow over the point-to-point link.

A caveat to direct encapsulation is the requirement that the end systems involved with the circuit must reside on Token Ring networks. Again, as with FST encapsulation, queuing is not as granular as TCP, thereby limiting the capability to prioritize the traffic between the peers. The up node to direct encapsulation is that it is fast-switched, enabling DLSw+ to process more packets per second than TCP encapsulation.

DLSw Lite

Encapsulating DLSw+ with LLC2 is known as DLSw Lite. It supports all the best features of DLSw+, using a four-byte LLC2 overhead with the 16-byte DLSw header. Currently, it is only supported over frame relay networks, assuming that the destination router is also the WAN router. The end systems, however, can be connected to Token Ring, SDLC, QLLC, or Ethernet networks.

Since DLSw Lite essentially uses a point-to-point connection, link failures disrupt the circuit, but recovery can be achieved through multiple peers at the destination location. DLSw Lite is processed-switched and therefore it does not achieve the packet-per-second capabilities of FST or direct encapsulation but meets the throughput of TCP encapsulation. Since DLSw Lite does not use TCP encapsulation, it cannot perform granular prioritization like that found with TCP encapsulation.

Encapsulation Overhead

As discussed, each encapsulation method incurs a different amount of overhead for each frame sent. Some encapsulation techniques, such as TCP and DLSw Lite, however, make up for their frame overhead by reducing the amount of unproductive traffic over the DLSw peer connections. Additionally, TCP encapsulation allows for multiple end system frames to be placed into a single TCP segment that will further reduce link overhead and provide more throughput to the end systems.

Cisco's Enhanced Modes of Operation

The Cisco DLSw+ implementation of DLSw RFC 1795 requires that the Cisco DLSw router be able to communicate with non-Cisco DLSw routers as well as with Cisco RSRB routers. In order for this to occur, the DLSw+ implementation provides three modes of operation:

- Dual
- Standard
- Enhanced

Dual Mode

In the dual mode of operation, a Cisco router communicates with both DLSw+ peers and RSRB peers. This is useful as a means for a migration path from RSRB to DLSw+. In dual mode, both RSRB and DLSw+ definitions and functions coexist on the router. A router operating in dual mode can communicate with another router in only one of the encapsulation techniques.

Standard Mode

DLSw+ operates in standard mode when it detects that the participating router is using RFC 1795. Although this mode does not afford the use of many of the advanced features of DLSw+, it does allow features that are local to the router. For instance, if the Cisco router is the mainframe attached DLSw+ peer, it can perform load balancing, local learning of end systems attached to ports on the router, the explorer firewall function, and media conversion.

Enhanced Mode

In enhanced mode, the DLSw+ router detects that the peer is also a DLSw+ router, making all the features supported under the DLSw+ specification available. These include:

- Peer groups
- Border peers
- On-demand peers
- Border peer caching
- Peer selection
- Backup peers
- Port lists
- FST
- Direct encapsulation.

Availability Configurations

DLSw+-enhanced features enable the network designer to build large-scale SRB networks with high availability. A high-availability network not only provides multiple means of access but is built with the intent to support alternate connectivity options for each important location and corporate service. DLSw+ builds on this using load balance, redundancy, and backup peers.

Load Balancing

High-availability SNA networks generally require multiple channel gateway connections to the SNA mainframe. These channel gateway connections are either IBM 3745 Communications Controllers or Cisco *Channel Interface Processor* (CIP)-connected routers. Load balancing is the function of connecting SNA sessions through the multiple channel gateways in a balanced manner.

Figure 9-5 diagrams a multirouter central site location with IBM 3745 and Cisco CIP channel gateways. Each gateway uses the same MAC address as the channel gateway address for connection to the mainframe. DLSw+ can cache up to four peer connections that provide connectivity to a destination MAC address and up to four local ports that can also be used to reach the end system MAC address. The illustration in Figure 9-5 demonstrates not only duplicate MAC address load-balancing to the mainframe but depicts a migration path from IBM 3745 to Cisco CIP channel gateways.

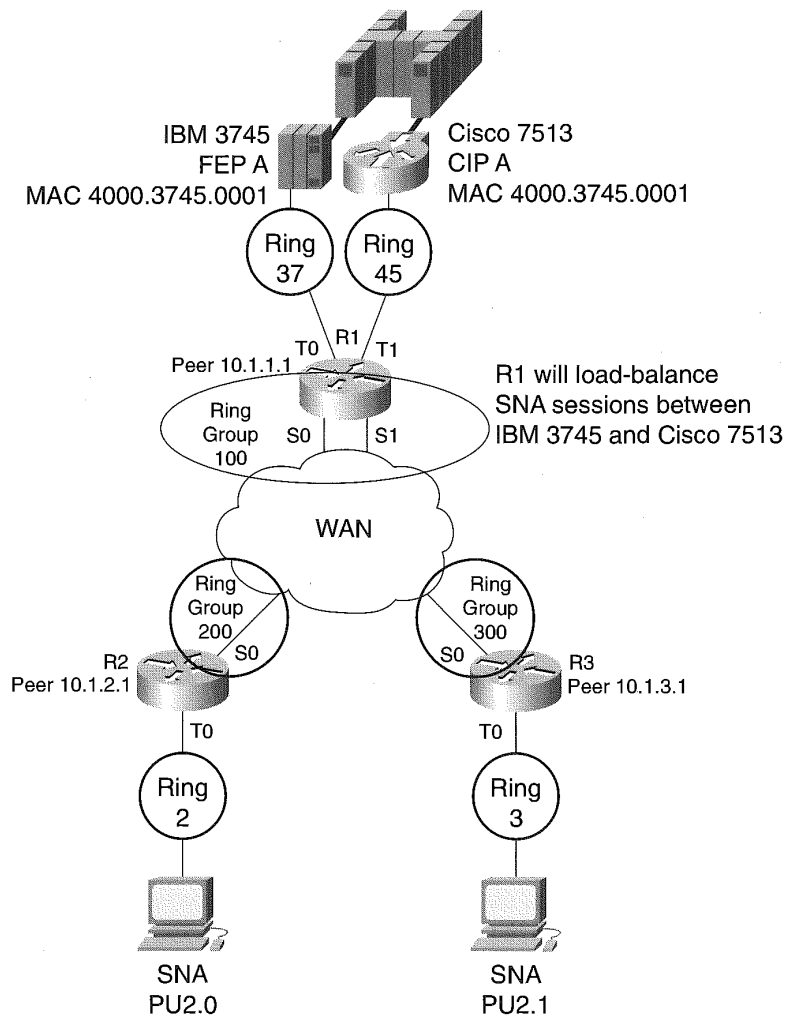
As mentioned earlier, if load balancing is specified for the DLSw+ definition, each new circuit set up will be established in a round-robin fashion over the next path in the cached sequence to the destination end system MAC address. Without specifying, load balancing DLSw+ will always select the first path in the cache as long as the path is available. The first path in the cache list is either the first peer to respond positively to the explorer frame, the peer with the least cost, or the local port over which the first positive response to the explorer frame was received.

Load balancing can be applied to all channel gateways, but they must be LAN-attached. The preferred LAN attachment is Token Ring, but Ethernet can be used.

Some caveats must be kept in mind when duplicating MAC addresses for channel gateways that are Ethernet-attached:

Figure 9-5

Load balancing between an IBM 3745 and Cisco 7513 channel-attached gateways



- A unique DLSw+ router and Ethernet bridge segment must be used for each duplicate MAC address.
- Load balancing is defined at the remote locations and not the central site routers.

These caveats are required to prevent loops that may occur when using Ethernet LANs. Loops are not a concern with Token Ring LANs since SRB inherently prevents loops.

Redundancy

Load balancing over a multiple-channel gateway is a form of redundancy. DLSw+, however, can also enable redundancy planning through the use of applying a cost to the DLSw+ connection. Placing a cost on a DLSw+ peer connection enables the connection to favor one peer over another to reach the same destination end system. Using a cost on a remote router also allows you to split SNA traffic from different LAN ports to different central site routers. Higher-cost peers can be used, should they receive a positive reply to the explorer frame before the lower-cost peer connection. This can be managed by setting a timer that forces DLSw+ to wait a period of time prior to selecting a peer. Another form of redundancy is the use of multiple peers and backup peers.


Backup Peers

Multiple active peers is the standard DLSw connectivity for peering as well as providing for alternate peering. This is not always prudent in design, however. The use of multiple active peers, for example, is not recommended when the alternate peer is at a disaster recovery location or at a secondary data center, or when two DLSw+ routers feed a single channel-attached router. These scenarios are better serviced using the backup peer feature of DLSw+. The backup peer feature requires the use of TCP or FST encapsulation. Figure 9-6 diagrams the use of backup peers.

In Figure 9-6, the mainframe is channel-attached to a Cisco CIP router. The CIP router connects via Token Ring to routers R1, R2, and R3. The remote routers R4 and R5 peer with routers R1 and R3 respectively. Router R2 is the backup router for both R4 and R5. If the connection between R4 and R1 were to fail, R1 would reestablish the SNA sessions to the mainframe through router R2.

Once connectivity is restored between R1 and R4, any new sessions to the mainframe are established over the DLSw+ connection between R1 and R4. Sessions already established through R2 remain active through R2 until they end or until a specified amount of time has expired, at which point the R2 router terminates the active remaining sessions.

Figure 9-6
Backup peer
topology using
DLSw+ backup peer
configuration



Using the backup peer feature in all remote routers reduces CPU and bandwidth requirements since only one DLSw+ connection is active at any given time. Additionally, a single backup router at the data center location can handle DLSw+ recovery services for hundreds of remote locations. The backup DLSw+ router at the data center does not perform any DLSw+ work until a primary DLSw+ data center router fails.

Performance Features

A recurring theme in router and bridge networking is the reduction or elimination of broadcasts. DLSw networks are no exception to this theme. DLSw+ addresses this issue through implementing broadcast domains with port lists, peer groups, border peers, on-demand peers, and dynamic peering.

Port Lists

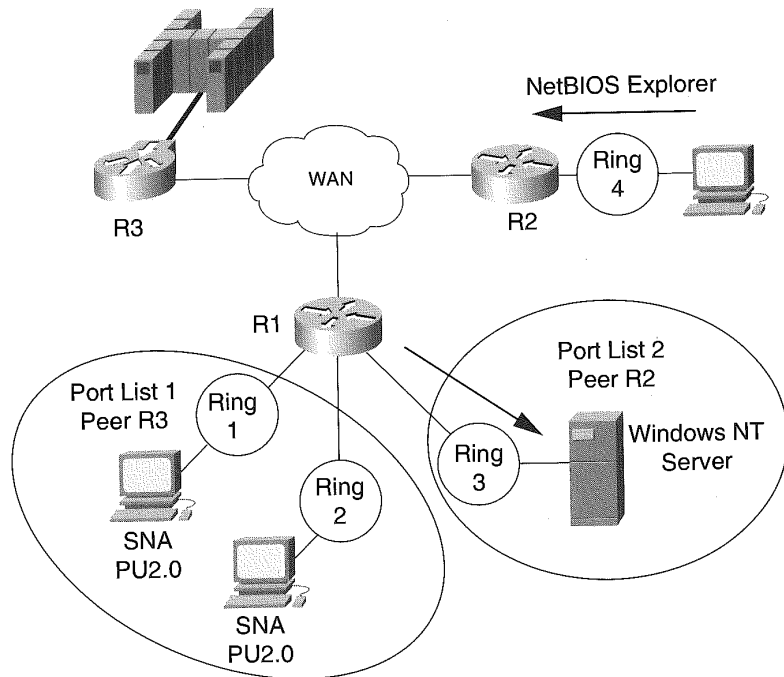
DLSw+ contains a feature for creating broadcast domains within the DLSw+ network. This is performed using the port list feature. Port lists enable the engineer to control the forwarding of broadcasts. They also allow differentiation between Token Rings and serial ports to establish the broadcast domains. All Ethernet ports are treated as a single domain due to the Ethernet bridge group definitions required. Figure 9-7 illustrates the use of port lists for establishing broadcast domains.

Router R1 in Figure 9-7 has defined Token Ring 1 and 2 as broadcast domains supporting SNA devices. Ring 3 on R1 is a broadcast domain of its own, supporting Windows NT servers using NetBIOS protocol. The workstation on R2 needs only to communicate with the Windows NT server on Ring 3 R1. The port list on R1 keeps all NetBIOS broadcasts from the workstation off Ring 1 and 2. The port lists can also stop all communications to any devices on Ring 1 and 2 from the workstation on R2.

Peer Groups with Border Peers and On-demand Peers

The DLSw standard calls for all DLSw peers to have a peer connection with all other DLSw peers. In a fully meshed network topology, this requires a calculated number of TCP connections expressed as

Figure 9-7
Establishing
broadcast domains
using DLSw+ port
lists



$$N \times (N-1)/2$$

The value for N is the number of routers used for DLSw. Therefore, in a network of 200 DLSw routers, there is a requirement for the network to establish 19,900 TCP connections in a fully meshed topology. This becomes a complex configuration and results in excessive explorer traffic. DLSw+ addresses the complexity and explorer explosion through the implementation of peer groups and border peers. By using the border peer feature, the number of TCP connections is reduced to $N-1$ and the number of explorer broadcasts are reduced by $N-2$ (this discounts the initial explorer from the source DLSw router).

In the 200 DLSw router network discussed above, for example, when using a single border peer in each of two peer groups, the active TCP connection requirement drops to 199 connections. By using border peers and two peer groups, a maximum of 198 explorers traverses the backbone for previously undiscovered end systems.

Understand that the full mesh topology discussed is logical and not necessarily physical. When a source and destination end system circuit is

established over a peer group connection using on-demand peers, it does not necessarily mean that the connection is a true point-to-point connection. In fact, the DLSw+ TCP packet travels to the established peer using the IP routing protocol algorithms. Therefore, it is important to ensure that ample bandwidth is available on the router connections that will form the actual physical path of the DLSw+ connection. Failing to recognize this need may lead to poor performance and the potential for low availability.

Dynamic Peers


Dynamic peers are useful in large networks that do not require the remote sites to have SRB connectivity to a central location all the time. Employing dynamic peers minimizes the number of central site routers due to the occasional use of the peer connections. Dynamic peers establish the DLSw+ peer connection only when an explorer arrives that meets a criteria matching the dynamic peer definition. The DLSw+ router dynamic peer definitions may have filter lists that apply SAP, MAC address, NetBIOS names, or even byte offset filters that will trigger the DLSw+ router to attempt a peer connection with the dynamic peer that meets the filter criteria. Since peer connections will send keepalive messages periodically to ensure that the remote peer is still active, using dynamic peers reduces CPU activity to service these unproductive messages.

Dynamic peers are also very useful for SNA dial-on-demand connectivity. Backup communication links for SNA locations can often require connectivity to a different location or a different router at the central site. Using dynamic peer definitions, DLSw+ automatically establishes a peer connection with the new peer and tears down the dialup connection after the last SNA session has become inactive.

CHAPTER

10

APPN Network Design



IBM developed Advanced Peer-to-Peer Networking (APPN) as a networking architecture in support of the small business SNA systems such as System/36, System/38, and the AS/400 computing platforms. The advantages of APPN architecture, however, were seen to provide an enhancement to traditional SNA networks that support large corporate mainframe environments. These enhancements include the following:

Peer-to-peer architecture: This eliminates the master/slave hierarchical architecture with its foundation in IBM's Advanced Communication Function/Virtual Telecommunications Access Method (VTAM) and Network Control Program (NCP).

Routing SNA traffic: APPN using Intermediate Session Routing (ISR) and the more advanced High Performance Routing (HPR) features facilitates the routing of SNA traffic that, prior to using APPN, had to be encapsulated in a routable protocol to be routed over a WAN. Using APPN, the SNA frames are not encapsulated. The HPR feature also enables non-disruptive rerouting of SNA traffic due to link or node outages. APPN-routing protocols also support the SNA Class of Service (COS) transport layer feature

Automatic Application Registration: The APPN architecture registers the location (the host containing the application) of all APPN applications on APPN network nodes (NNs). This eliminates network address or host name knowledge requirements when one application needs to communicate with another.

Full APPN topology database: APPN network nodes communicate with each other and build a complete network topology. The network topology database enables the APPN-routing algorithms to determine the least-cost path between the two end systems.


Support for non-APPN-capable SNA devices: APPN specifications provide for traditional SNA legacy devices and resources to access an IBM mainframe application over an APPN WAN. These device types, SNA physical unit type 2 (PU Type 2) resources, require IBM's VTAM for activation and session setup with a mainframe application. This preserves a corporation's investment in these resources while taking advantage of the APPN architecture.

APPN consists of four node types: network node (NN), end node (EN), low-entry networking node (LEN), and border node (BN). Each node has a

control point (CP), which is responsible for connectivity, session establishment, and communicating capabilities. Figure 10-1 diagrams an APPN network. The following lists the functions of each node type:

- *Network Node (NN)*: The APPN NN is the keeper of all APPN information. An NN maintains the network topology and directory databases, performs the route decision process, and then routes the traffic through the network. The CP of an NN establishes a CP-CP session with the CP of adjacent NNs and ENs. The NN CP manages and is the repository of its own resources, registered resources from attached end nodes, and entered resources of LENs. APPN nodes currently communicate over single-link transmission groups (TGs). NNs learn the network topology through dynamic topology updates exchanged over the CP-CP sessions with adjacent NNs on active TGs.
- *End nodes (EN)*: APPN ENs and their CP manage only their local resources. An EN can register its resources with an attached NN over a CP-CP session. An EN CP, on behalf of a local application resource, will request directory services of an attached NN and request routing information for a destination APPN resource.
- *Low-entry networking node (LEN)*: LEN nodes appeared in the first incarnation of APPN for use on System/36 and System/38 computing platforms. The LEN possesses a CP, but it does not participate in a CP-CP session with an adjacent node. The LEN CP does manage the local resources, but any peer-to-peer communications outside of the LEN must be manually defined since the LEN CP does not establish CP-CP sessions with other APPN nodes. Likewise, LEN node resources must be predefined on any other APPN node that will require peer-to-peer communications with the LEN node application resources.
- *Border node (BN)*: A BN enables the subnetworking of large APPN topologies into smaller, more manageable networks. A BN will connect to multiple APPN networks as either an EN representing the owner of all the APPN LUs registered in the APPN network or as an NN for another APPN network managing routing tables and directory services between the multiple networks. Cisco routers do not at this time support the BN function. It is only supported on AS/400 and mainframes, executing IBM VTAM V4R4.

Figure 10-1
An APPN node
topology.



Cisco Support of APPN

The Cisco IOS software supports NN functionality. This enables the Cisco router to do what it does best, routing information over a WAN backbone. Cisco has licensed the APPN code from IBM for porting to Cisco IOS to ensure full functionality and interoperability with non-Cisco NNs. Because

APPN is an upper-layer protocol, it can use lower-layer functions for transport between APPN peer resources. As such, Cisco IOS software supports the following mechanisms for transporting APPN over a Cisco router-based backbone network:

- APPN support over pure IP router backbone networks is possible using DLSw+ or RSRB encapsulation and network transport.
- Use of the connection network concept is possible over DLSw+, RSRB, Frame Relay, and ATM-based networks, as well as Token Ring and Ethernet LANs.
- Cisco APPN NN support on a channel-attached Cisco router allows VTAM to directly interface with the network as either an end node or a network node.
- Cisco routers can act as an APPN Dependent Logical Unit Requestor (DLUR) to VTAM's Dependent Logical Unit Server (DLUS) functions, thereby enabling legacy non-APPN capable SNA physical units to communicate to the mainframe over APPN.
- Cisco router APPN network nodes support both Intermediate Session Routing (ISR) and High-Performance Routing (HPR).
- Native Service Point (NSP) support for SNA network management to IBM's NetView or Sterling Software's Net/Master, along with SNMP network management support.

In addition to the above-mentioned features, APPN can be transported using several options, providing a robust solution for using APPN in a Cisco router-based network. The options are as follows:

- DLSw+ over Frame Relay
- BNN/BAN for Frame Relay
- Frame Relay host
- Fast Sequenced Transport (FST)
- Direct encapsulation

Deciding Factors on Using APPN

Several factors contribute to choosing APPN as the network protocol for carrying SNA traffic over a multiprotocol backbone. Among these are support for mission-specific applications riding on SNA protocol, the requirement of meeting service levels on these applications with guaranteed

delivery, the need for office-to-office SNA communication requirements, and support of legacy SNA devices over a multiprotocol network.

APPN deployment must be considered carefully. APPN functionality should be used when considering the following network criteria:

- Granular prioritization is in effect, down to the application running on an end user station.
- The majority of the network traffic is SNA-based.
- The network infrastructure and skill set is primarily SNA.
- LU6.2 applications are being deployed for SNA client/server business solutions between mainframes, AS/400s, and SNA emulated on UNIX servers.
- Peer-to-peer SNA routing is required between SNA hosts without mainframe connectivity, such as AS/400-AS/400 connectivity over the common enterprise-wide network infrastructure.
- VTAM V4R1 or higher is being implemented or is in use.

The Dominant Protocol Is SNA

Although TCP/IP applications have been embraced for providing corporate mission-critical applications, it is still estimated that 80 percent of all applications require SNA-derived data. SNA, in many organizations, still provides the bulk of applications for running the business processes. Analysis of the application characteristics at the branch office and corporate networks will provide a direction as to whether APPN is applicable to the network. In general, if the application base is using APPN LU6.2 across the WAN, then implementing APPN on backbone routers is a tactical solution for ensuring service levels of SNA integrated with a multiprotocol backbone network.

Class of Service (COS)

Class of Service (COS) along with transmission priority (TP) are SNA functions for prioritizing SNA traffic based on application service-level requirements. This is compared to custom queuing, which prioritizes at the end system level. COS and TP prioritize the traffic for multiple applications that may occur from a single end system.

In typical SNA networks, COS is focused on VTAM to NCP connections, NCP-NCP connections, and NCP-PU T2/2.1 connections. This is illustrated in Figure 10-2. Here we see a typical SNA network with and without routers. Note that when routers are added to the network they will fall in line with non-router topology, thereby requiring the COS function for continued delivery of service levels.

Office-to-Office Connectivity

Assigning network node functionality in a Cisco router located at a branch office is dependent on the needs of the branch office. Figure 10-3 illustrates the two types of functional requirements. If the branch office requires communication with only the data center, then connecting DSLw+ to either the backbone router, which serves as an NN, or to the data center router using APPN NN functionality is the appropriate solution.

If there is an ongoing requirement that the branch office communicate with resources in other branch offices, however, then it may be cost-justifiable to enable NN functionality on the branch office routers involved, as long as they have a direct WAN link. With a direct link between branch offices, the branch office router will route the APPN traffic over the direct link, rather than the backbone network node. Placing an APPN NN router at the branch office enables SNA COS functionality from the branch office to the data center through the entire Cisco WAN. This ensures service levels for applications requiring guaranteed performance criteria.

Figure 10-2
COS points in a typical SNA network.


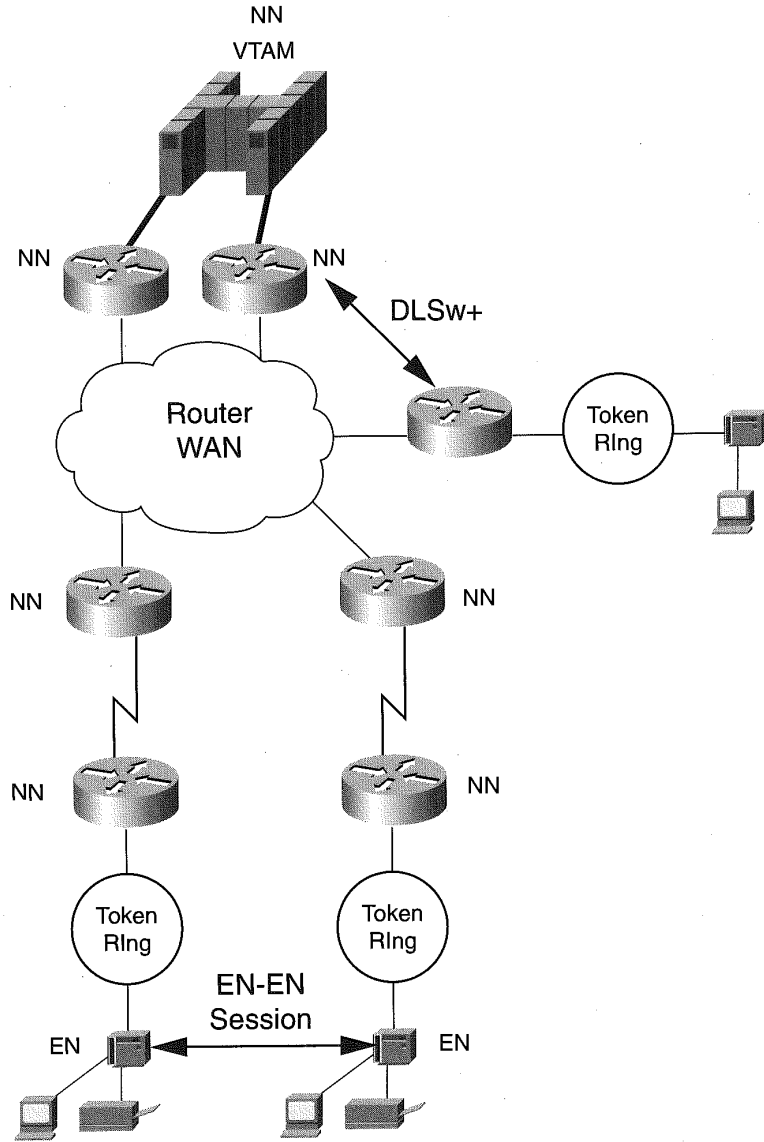


Figure 10-3
Branch office APPN functional solutions.



Non-APPN-capable SNA Resources

Support for legacy, non-APPN-capable SNA resources has not been available until recently with APPN. Many corporate SNA networks still have IBM 3274 and IBM 3714 cluster control units supporting IBM 3270 applications. These control units and their resources require the traditional hierarchical master/slave architecture of SNA to communicate with the SNA applications residing on the mainframe. In networks where these devices are still in use but APPN is the chosen protocol for connectivity, a new feature named Dependent Logic Unit Requestor/Dependent Logical Unit Server (DLUR/DLUS) is employed.

Dependent Logical Unit Requester/Server


Support for transporting legacy SNA traffic over APPN is implemented on Cisco routers using Dependent LU Requestor (DLUR) in concert with VTAM V4R2 or higher implementing the Dependent LU Server (DLUS) function. The support on VTAM V4R2 or higher is also referred to as Session Services Extension. Using DLUS and DLUR legacy SNA devices attached to a router implementing DLUR can have their session and data transported over an APPN HPR network, thereby gaining all the advantages of a dynamic transport. Figure 10-4 diagrams a DLUR/DLUS connection supporting legacy SNA PU devices. Note that there is no need for the legacy SNA PU to have an owning VTAM (SSCP-PU) session with which to enable LU-LU sessions. This function is emulated by the DLUR feature defined on the Cisco router.

Network Node Placement

In general, if a routing decision needs to be made at a location, then the router at that location may be a good candidate for becoming an APPN NN. For example, Figure 10-5 shows a network configuration containing branch offices and a data center. In this configuration, we see that some of the branch offices have connectivity to the backbone with a single connection. In branch office locations where a single link connects the location to the

Figure 10-4

APPN support for legacy SNA devices using DLUR/DLUS.



backbone and application prioritization is a requirement, placing an APPN NN at the branch office will afford COS prioritization for the applications at the branch office.


In order to implement APPN NN on Cisco routers, however, there may be a substantial increase in cost due to memory and hardware requirements for effectively running APPN NN functions on the routers. If the cost becomes prohibitive, then the network design may be better served by using DLSw+ to deliver the SNA traffic to an APPN NN residing in the backbone. Once the data is in the backbone, it will be prioritized for delivery to the mainframe.

Many installations utilize APPN NN functionality on both backbone and data center channel-attached routers. Although this configuration works in some instances, the use of APPN NN functions on channel-attached routers along with WAN services may put undue stress on the router processor and cause scalability problems. A solution to enable channel-attached routers to concentrate solely on the delivery of APPN data is to apply the three-layer hierarchical model at the data center. Figure 10-6 illustrates this design. The channel-attached routers perform as APPN NNs to the mainframe and connect to the distribution routers over an ATM backbone. The distribution routers then provide APPN NN and WAN support to the rest of the network.

Performance Considerations

APPN is essentially a link-state routing protocol. Topology updates are used between NNs when any NN becomes cognizant of a new link, new resource, or the loss of a link or resource. Because of this, the size of the APPN network and the number of network nodes results in performance concerns. The larger the network, the more network traffic, high memory, and CPU utilization for maintaining a large topology database and COS tables. In

Figure 10-5
 Branch office APPN
 NN scenarios for
 supporting NN
 functionality on
 branch office routers.



addition, APPN ENs request a search for connecting to a partner application. The search is known as an APPN LOCATE and itself may end up consuming bandwidth and processor utilization. The goal of defining a well-behaved scalable APPN network is to control the number of TDUs and minimize LOCATE search requests.

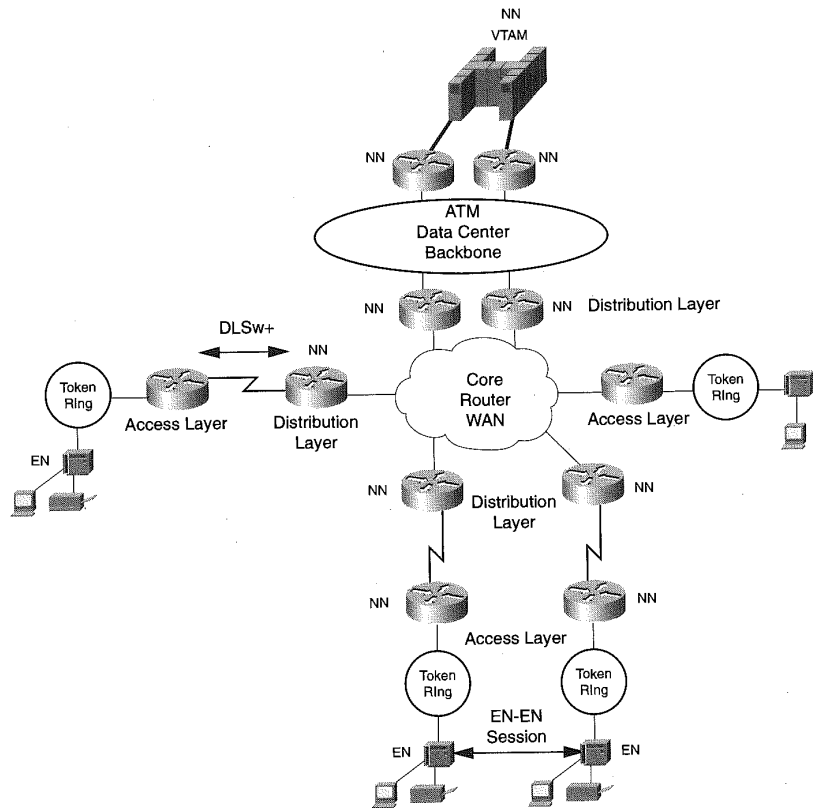
Topology Database Updates

The topology database update (TDU) message provides several different pieces of information to an NN. This information includes

- Sending node characteristics
- Node and link characteristics of other network resources
- The most recent update sequence number, or the resources being described

Figure 10-6

Applying the three-layered hierarchical routing model to APPN for data center connectivity.



When an NN learns of a topology change, it sends TDUs over every active link that attaches an adjacent NN. This flooding technique is used for rapid convergence. The more NNs in a network, the more TDUs are flowing, and therefore the greater the risk of high consumption of CPU cycles, memory, and bandwidth. A goal for providing a scalable and optimal performing network is one that uses a minimum number of links, network nodes, and hence a minimal number of CP-CP sessions.

Connection Network

The number of links used for establishing an APPN network depends on the media used for the connections. An APPN network will view the connections of nodes over a LAN or an ATM LANE network, or Cisco-based

DLSw+ or RSRB networks are seen as a large fully meshed APPN network since these encapsulation techniques emulate LANs. This topology is shown in Figure 10-7.

As in DLSw, a fully meshed network will contain $N \times (N-1) / 2$ CP-CP sessions. For example, an APPN network containing 50 network nodes will have 1,225 CP-CP sessions in a full mesh topology. This can cause extensive TDU flows and jeopardize the scalability and performance of the network.

Since TDUs flow on the transmission group (TG) links that connect adjacent NNs, reducing the number of links will reduce the number of TDUs. In a LAN environment, this is accomplished using the concept of a connection network. A connection network is the definition of a virtual routing node (VRN) that is defined as an NN to all of the NNs on the LAN.

In an APPN network, an NN can be defined as an NN Server (NNS). A NNS maintains all the topology information of all adjacent nodes on a LAN. Each NN in a connection network configuration defines a link only for connecting to the NNS and to the VRN. When an NN establishes its link with the NNS, the NNS sends its topology database to the NN. The NN sends TDUs to the NNS, defining its own characteristics, its link to the NNS, and its link to the same VRN. Once all the other NNs have established a link with the NNS, the connection network is fully functional. When an NN on the connection network requests services of the NNS for establishing a session between resources that exists off of NNs attached to the connection network, it notifies the requesting NN that the destination is on the same network and that a direct route is to be used.

Using the concept of connection networks on LANs drastically reduces the number of actual links defined in the network, thereby improving scalability and performance while providing any-to-any connectivity.

Number of CP-CP Sessions

Reducing TDU flows over media that cannot support the connection network concept is another design challenge for a scalable APPN network. Since NNs exchange TDUs over the CP-CP sessions, reducing the number of CP-CP sessions over an APPN WAN reduces the number of TDU flows. Designing a WAN APPN topology in a fully meshed WAN configuration calls for only two CP-CP sessions for each NN.

Shown in Figure 10-8, a fully meshed link environment allows for two CP-CP sessions between the adjacent NNs in the figure. In such an environment, the loss of a single link to either adjacent node still enables TDUs to flow over the remaining link, thereby allowing routing and updates to flow.

DLSw+ or RSRB networks are seen as a large fully meshed APPN network since these encapsulation techniques emulate LANs. This topology is shown in Figure 10-7.

As in DLSw, a fully meshed network will contain $N \times (N-1) / 2$ CP-CP sessions. For example, an APPN network containing 50 network nodes will have 1,225 CP-CP sessions in a full mesh topology. This can cause extensive TDU flows and jeopardize the scalability and performance of the network.

Since TDUs flow on the transmission group (TG) links that connect adjacent NNs, reducing the number of links will reduce the number of TDUs. In a LAN environment, this is accomplished using the concept of a connection network. A connection network is the definition of a virtual routing node (VRN) that is defined as an NN to all of the NNs on the LAN.

In an APPN network, an NN can be defined as an NN Server (NNS). A NNS maintains all the topology information of all adjacent nodes on a LAN. Each NN in a connection network configuration defines a link only for connecting to the NNS and to the VRN. When an NN establishes its link with the NNS, the NNS sends its topology database to the NN. The NN sends TDUs to the NNS, defining its own characteristics, its link to the NNS, and its link to the same VRN. Once all the other NNs have established a link with the NNS, the connection network is fully functional. When an NN on the connection network requests services of the NNS for establishing a session between resources that exists off of NNs attached to the connection network, it notifies the requesting NN that the destination is on the same network and that a direct route is to be used.

Using the concept of connection networks on LANs drastically reduces the number of actual links defined in the network, thereby improving scalability and performance while providing any-to-any connectivity.

Number of CP-CP Sessions

Reducing TDU flows over media that cannot support the connection network concept is another design challenge for a scalable APPN network. Since NNs exchange TDUs over the CP-CP sessions, reducing the number of CP-CP sessions over an APPN WAN reduces the number of TDU flows. Designing a WAN APPN topology in a fully meshed WAN configuration calls for only two CP-CP sessions for each NN.

Shown in Figure 10-8, a fully meshed link environment allows for two CP-CP sessions between the adjacent NNs in the figure. In such an environment, the loss of a single link to either adjacent node still enables TDUs to flow over the remaining link, thereby allowing routing and updates to flow.

Figure 10-7

Reducing TDU flows by using connection networks.

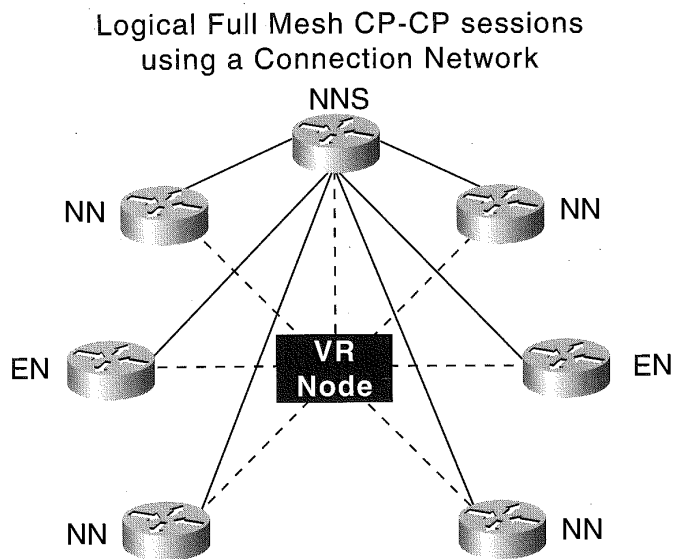




Figure 10-8
Using dual CP-CP sessions for reducing TDUs.

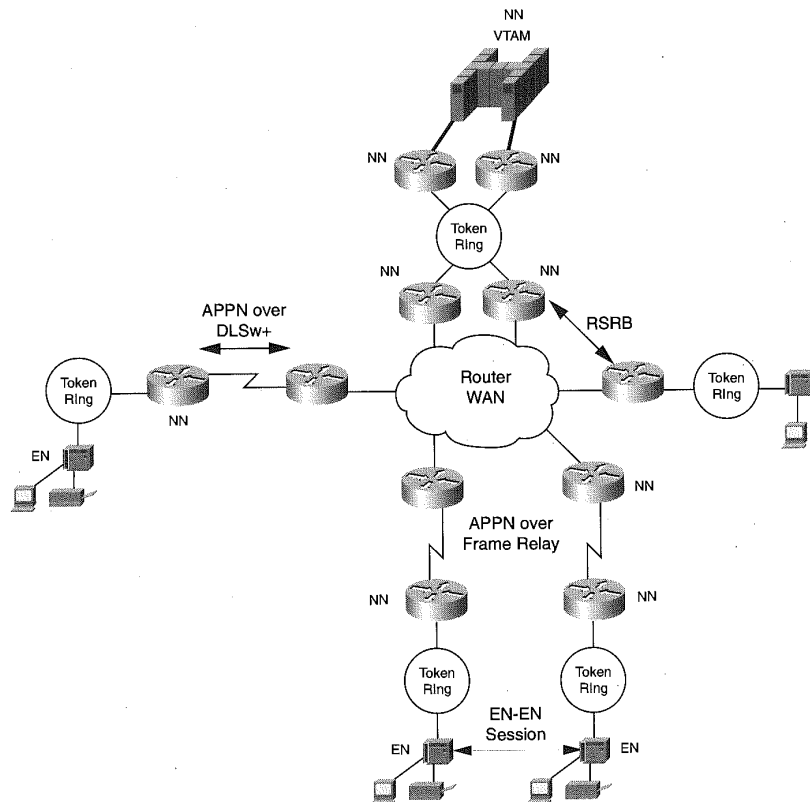


Number of Network Nodes

At first glance, it seems logical that all routers within the Cisco router network handling SNA traffic should implement APPN NN functionality. APPN NNs use TDU messages for updating adjacent NNs of topology changes. Placing the APPN NN function on all the routers could then lead to TDU congestion and therefore network scalability issues. It is more prudent to keep the number of APPN network nodes to a minimum.

A methodology for reducing the number of NNs in an APPN network is by pushing the NNs out to the edges of the network and relying on other transport mechanisms to deliver the APPN traffic between the edge NNs. Figure 10-9 illustrates this topology. Cisco offers DLSw+, Frame Relay Access Server (FRAS) with Boundary Network Node (BNN), Boundary Access Node (BAN), or RSRB.

Figure 10-9
Reduced NN
topology.



Using DLSw+ allows a network engineer to apply prioritization at the edge router for outbound packets, but the prioritization is lost in the intermediate routers used for connecting the edge NNs. The DLSw+-enhanced features along with peer groups, border peers, backup peers, and load balancing provide the means for designing a highly scalable APPN network.

If the backbone connectivity is using Frame Relay, Cisco FRAS BNN/BAN can be used for connecting the edge APPN nodes. In this type of configuration, the data center router performs both NN and frame relay WAN functionality for delivering the APPN traffic. Using frame relay in such a situation provides an effective scalable solution, but it does lose the COS prioritization of traffic over the backbone network.

Cisco RSRB can also be used for connecting the edge APPN nodes. Using RSRB provides non-disruptive rerouting due to link failures, as does DLSw+. RSRB is being phased out by Cisco, however, and is not a recommended design solution for reducing the number of NNs.

LOCATE Search

APPN end nodes request assistance from their network node server (NNS) for locating a network resource during application-to-application establishment. If the location of the requested resource is unknown to the NNS server, a LOCATE search request is sent over all the CP-CP sessions maintained with other network nodes by the NNS on behalf of the requesting end node. This dynamic search in large networks can contribute to unforeseen bandwidth and processing utilization. Reducing the number of LOCATE search requests flowing through the network contributes to the successful scaling of large APPN networks. The following techniques and features of Cisco's APPN implementation can be used to reduce the LOCATE search flows:

- Use of a Trivial File Transfer Protocol (TFTP) server for recovering the last-known directory cache of a Cisco APPN network node using the Safe-Store of Directory Cache function. A restarted Cisco network node obtains the resource directory from the TFTP server, rather than from broadcasting LOCATE requests throughout the network.
- Using partial directory entries coded within the Cisco router APPN configuration identifying the full resource name or a wildcard partial name, along with the name of the owning end node or network node.

Using DLSw+ allows a network engineer to apply prioritization at the edge router for outbound packets, but the prioritization is lost in the intermediate routers used for connecting the edge NNs. The DLSw+-enhanced features along with peer groups, border peers, backup peers, and load balancing provide the means for designing a highly scalable APPN network.

If the backbone connectivity is using Frame Relay, Cisco FRAS BNN/BAN can be used for connecting the edge APPN nodes. In this type of configuration, the data center router performs both NN and frame relay WAN functionality for delivering the APPN traffic. Using frame relay in such a situation provides an effective scalable solution, but it does lose the COS prioritization of traffic over the backbone network.

Cisco RSRB can also be used for connecting the edge APPN nodes. Using RSRB provides non-disruptive rerouting due to link failures, as does DLSw+. RSRB is being phased out by Cisco, however, and is not a recommended design solution for reducing the number of NNs.

LOCATE Search

APPN end nodes request assistance from their network node server (NNS) for locating a network resource during application-to-application establishment. If the location of the requested resource is unknown to the NNS server, a LOCATE search request is sent over all the CP-CP sessions maintained with other network nodes by the NNS on behalf of the requesting end node. This dynamic search in large networks can contribute to unforeseen bandwidth and processing utilization. Reducing the number of LOCATE search requests flowing through the network contributes to the successful scaling of large APPN networks. The following techniques and features of Cisco's APPN implementation can be used to reduce the LOCATE search flows:

- Use of a Trivial File Transfer Protocol (TFTP) server for recovering the last-known directory cache of a Cisco APPN network node using the Safe-Store of Directory Cache function. A restarted Cisco network node obtains the resource directory from the TFTP server, rather than from broadcasting LOCATE requests throughout the network.
- Using partial directory entries coded within the Cisco router APPN configuration identifying the full resource name or a wildcard partial name, along with the name of the owning end node or network node.

- Defining VTAM as the Central Directory Server (CDS) and employing the CDS/Client function on Cisco APPN network nodes to query the CDS on VTAM for the location of a resource.
- End node resource registration requests NN to register the resource with the CDS. This quickly populates the CDS and adds to the reduction of LOCATE search flows.

Combining these techniques with those discussed for reducing TDU flows will greatly enhance stability, performance, and scaling of the APPN network.

Recovery Techniques

As with all network designs, link recovery is paramount to providing high availability. The most proven techniques for an APPN network are a secondary WAN backup link, dual active WAN links and routers, and APPN DLUR backup using Cisco CIP router connectivity to VTAM on the mainframe.

Secondary WAN Link

Typically, a secondary WAN link is a switched (dial-up) connection. It is usually an ISDN BRI connection at 56/64 Kbps or 128 Kbps. Using Cisco IOS software features, a secondary ISDN connection can lay dormant and be activated on demand to the adjacent APPN node, should the primary link fail. Although the auto-activation of the ISDN line is possible, deactivation must be done manually.

Dual Routers/Dual WAN Links

Using two routers and two active links from the remote node to the central location enables active redundancy. In such a configuration, should either link (TG) become inoperative, the end-to-end connection continues to flow uninterrupted, with HPR as the routing protocol used on the router NNs.


High-Performance Routing

High-performance routing is APPN's second generation of routing protocol. Using HPR SNA APPN sessions can be rerouted over the backbone network without disruption to the SNA session. HPR support includes the following:

- Reduced processing at intermediate nodes by moving the error processing, flow control, and segmentation to end node
- Non-disruptive rerouting of data due to path failures
- Use of Adaptive Rate-Based (ARB) flow control, providing an optimal method for high-bandwidth environments

HPR support on Cisco routers is found in IOS V11.3. Figure 10-10 illustrates the use of Cisco routers in an APPN HPR network.

Figure 10-10
APPN HPR with Cisco routers.



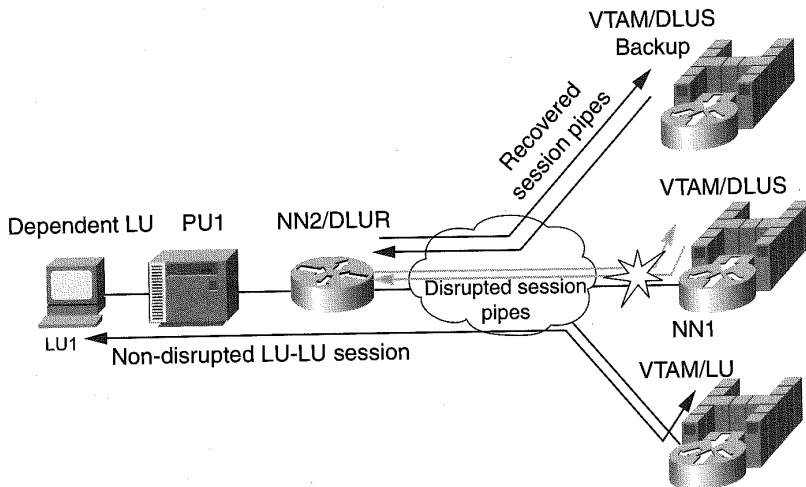
SSCP Takeover of DLUR

DLUR/DLUS is an APPN function that supports legacy SNA connectivity. VTAM enables a process called takeover/giveback. These functions together enable APPN DLUR legacy connections to continue and have their sessions owned by a backup VTAM if the primary VTAM were to fail. The takeover and giveback process is initiated by VTAM commands.

Takeover is the capability of a VTAM System Services Control Point (SSCP) to take ownership of a dependent device if the original SSCP should fail. Giveback is the return of ownership back to the original owning SSCP. Both takeover and giveback are non-disruptive to all established cross-domain sessions. DLUS supports a similar functionality, as shown in Figure 10-11.

The initial DLUS session pipes between the DLUR and the DLUS have failed. A backup DLUS VTAM activates new session pipes to the DLUR of the dependent resources. Once the activation sequences have been accomplished, new SSCP-PU and SSCP-LU sessions are established, and subsequent LU-LU session requests are satisfied through the new DLUR-DLUS session pipes. The takeover and giveback of the DLUR-DLUS session pipes are non-disruptive to current LU-LU sessions.

Figure 10-11
DLUR/DLUS support
for VTAM ANS and
takeover/giveback
functions.



Queuing and Prioritization

Managing the sensitivity of APPN connections to network delay is paramount in providing high availability. Traffic priorities and the reservation of bandwidth can assist in protecting the service-level requirements of APPN traffic. The Cisco IOS software features for the queuing and prioritization of traffic on Cisco routers also apply to APPN. For more information on these methods consult Chapter 9, "DLSw+ Network Design."

APPN Buffer and Memory Management

Because APPN network nodes store a network topology database along with a directory database and COS tables, a considerable amount of memory is consumed on the router. The Cisco IOS software allows the specific reservation of the memory used by the APPN function on the router. The memory on the router used by APPN can be specified as dedicated or shared.

Specifying the memory as shared is accomplished by defining to the router the maximum amount of memory the APPN function can use for the databases and tables. If a minimum amount of memory is specified instead, the router will dedicate the value specified and reserve that amount of memory for APPN functionality only. Specifying a minimum amount of memory should be used with caution, as other router processes will not be able to use the reserved amount of memory. The amount of memory defined for APPN use is a determining factor in the number of supported sessions.

Each APPN session traversing the router uses buffers allocated in the APPN memory space. The buffers regulate the inbound and outbound traffic through the router. The APPN buffer space is defined as a percentage of the overall APPN memory allocation. During buffer memory shortages, the APPN function employs statistical flow control mechanisms to avoid severe buffer congestion or deadlock conditions to the buffer shortage. Table 10-1 lists the estimates for APPN memory allocation on a Cisco router.

Executing APPN on a Cisco 2500 series router requires a minimum of 8 MB of DRAM and 8 MB of flash memory. The limiting factors in this environment are the number of sessions supported by the CPU and the number of interfaces. The enterprise plus APPN subset requires 16 MB of flash memory.

Table 10-1

Estimating Table for
APPN Memory
Requirements

| APPN Memory Requirements | KB |
|---|----------------|
| Starting APPN process (Control blocks, stacks, etc.) | 760.0 |
| Each APPN port defined in the router configuration | 1.5 |
| Each local APPN link, defined but not active in the router configuration | 1.5 |
| Each active APPN/LEN/subarea link that does not have CP-to-CP sessions in the router configuration | 6.5 |
| Each APPN link that has CP-to-CP sessions in the router configuration | 12.5 |
| Each connection network defined in the router configuration | 1.5 |
| Each network node in the network topology that is not adjacent to this node in the router configuration | 0.4 |
| Each connection between network nodes in the network topology in the router configuration (do not include connections of which this node is an end point) | 0.9 |
| Each PU served by DLUR on this node in the router configuration | 1.5 |
| Each LU served by DLUR on this node in the router configuration | 0.9 |
| Each DLUS with which this node has a DLUR/S connection active in the router configuration | 15.0 |
| Each active intermediate session in the router configuration | 1.3 |
| Each cached or registered resource in the APPN directory in the router configuration | 0.3 |
| Total static APPN memory requirements in the router configuration | KB sum |
| Total operational APPN memory requirement without buffers in the router configuration | 2*(KB sum) |
| Total operational APPN memory requirement with buffers in the router configuration | 2*(2*(KB sum)) |


The Cisco 4000 and 4000 M series require a minimum of 16 MB of DRAM, 4 MB of flash memory, and 4 MB of shared DRAM. If using the Cisco 4500, most configurations function appropriately with 32 MB DRAM, 4 MB of flash memory, and 16 MB of shared DRAM. The Cisco 4500-M or 4700-M requires a minimum of 16 MB of DRAM, 4 MB of flash memory, and 16 MB of shared DRAM. Larger networks supported by the 4500-M or 4700-M series routers require a minimum of 32 MB of DRAM.

The Cisco 7000 and 7500 series routers require a minimum of 16 MB of DRAM for very small networks and 64 or 128 MB of DRAM for large APPN networks.

CHAPTER

11

ISDN and DDR Design



Integrated Services Digital Network (ISDN) was created with the idea of being able to digitize all forms of communication for transport over telephone networks. Router implementation of ISDN has been in the areas of providing temporary extra bandwidth, temporary casual remote access, and as a high-speed switched backup connection for recovering remote locations after a dedicated line outage.

Dial-on-demand routing (DDR) enables Cisco routers to communicate over temporary public-switched telephone networks (PSTN) for on-demand services. DDR is functional using any of the switched interface connections available to a Cisco router platform. DDR can be used for providing a switched backup connection to support outages in a dedicated WAN link topology or it can be used as a means of providing temporary connectivity based on "interesting" packets that flow through the router interfaces.

Site Options

The Cisco Systems products offer many varied solutions for both remote and centralized ISDN connections. The more popular Cisco router ISDN solutions are listed in Table 11-1.

Table 11-1

Sample ISDN Offerings on Cisco Router Platforms

| Cisco Router Platform | ISDN Interfaces |
|-------------------------------------|---|
| Cisco 700 series | One BRI |
| AccessPro PC card | One BRI |
| Cisco 1003/4 | One BRI |
| Cisco 2503/4 | One BRI |
| Cisco 2516/17 | One BRI |
| Cisco 3600 series router (per NPM) | Four or eight BRI's per NPM, one or two PRI per NPM |
| Cisco 4000 series router (per NPM) | Four or eight BRI's per NPM, one PRI per NPM |
| Cisco 7000 series/Cisco 7500 series | One or two PRI's per MIP |

For a complete description of these platforms and the ISDN offerings on each, consult Chapter 2, "Cisco Router Hardware," or the appropriate Cisco systems product catalog. Of the Cisco router platforms listed, the 3600, 4000, and 7000 series routers are traditionally used at the central site, which acts as the hub of a larger hub and spoke topology, offering one or two ISDN PRI lines for connection to the dialer cloud.

Central Office Switch Considerations

In preparing for designing the ISDN network, it is imperative that the capabilities of the provider central office switches be discovered early on in the design effort. Telephone companies may use a wide variety of ISDN switches that are based on the standard but have been modified by the manufacturer. Knowing the type of switch being used by the telephone company for central and remote connections will ease your implementation process.

The predominant ISDN switches used in North America are the AT&T 5ESS and the Nortel DMS-100. Prior to the release of the National ISDN-1 software standard for ISDN, calls from an AT&T ISDN phone would fail when connecting to a Nortel DMS-100. Likewise, a Nortel ISDN phone set would fail on connecting to an AT&T 5ESS ISDN switch. Both the AT&T 5ESS and Nortel DMNS-100 ISDN switch support up to 100,000 local loop connections. The switches are geared towards either PRI or BRI line configurations. Table 11-2 lists the various North American switches and their keyword definitions within the Cisco IOS router definitions.

ISDN networks are now available around the world. In countries outside of North America, however, other ISDN switches may be deployed in the central offices that are not AT&T 5ESS or Nortel DMS-100. The possible switch types used on other countries of the world are listed in Table 11-3.

In some cases, the switch being used deactivates the Layer 2 functions of the D channel on a BRI connection when no calls are active. In this case, the router must perform a Terminal Equipment Identification (TEI) negotiation when a call is first made, rather than when a router powers-up and activates the BRI interface.

Table 11-2

AT&T 5ESS and Nortel DMS-100 Switch Specification for Router Definitions

| ISDN Switch Type | Router Keyword |
|----------------------------------|----------------|
| AT&T basic rate switch | Basic-5ESS |
| Nortel DMS-100 basic rate switch | Basic-DMS100 |
| National ISDN-1 switch | Basic-NI1 |
| AT&T 4ESS (ISDN PRI only) | Primary-4ESS |
| AT&T 5ESS (ISDN PRI only) | Primary-5ESS |
| Nortel DMS-100 (ISDN PRI only) | Primary-DMS100 |

Table 11-3

ISDN Switch Types Used in Other Countries and the Router Configuration Keyword

| Country | Switch Type | Router Keyword |
|----------------|------------------------------------|----------------------------|
| Australia | TS013 BRI switch, TS014 PRI switch | Basic-TS013, primary-TS014 |
| France | VN2 ISDN switch, VN3 ISDN switch | VN2, VN3 |
| Germany | 1TR6 ISDN switch | Basic-1TR6 |
| Japan | NTT ISDN switch, ISDN PRI switch | NTT, primary-NTT |
| New Zealand | NET3 ISDN switch | Basic-NZNET3 |
| Norway | NET3 ISDN switch (phase one only) | Basic-NWNET3 |
| United Kingdom | NET3 BRI switch, NET5 PRI switch | Basic-NET3, primary-NET5 |

ISDN connections are established through the use of service profile identifiers (SPIDs). SPIDs are assigned by the ISDN provider to identify the line configuration of the BRI service. Using SPIDs, voice, data, video, and fax can share the local loop. The DMS-100 and National ISDN-1-supported switches require a SPID. The AT&T 5ESS does not always require a SPID, but specifying a SPID will not be detrimental as long as the carrier provides the correct corresponding SPID.

An AT&T 5ESS switch can support up to eight SPIDs per BRI, while the DMS-100 and National ISDN-1 switches support only up to two SPIDs per BRI. The capability of the AT&T 5ESS to support multiple SPIDs per BRI B channel allows multiple services simultaneously over the same B channel. The DMS-100 and National-ISDN-1 allows only one SPID per BRI channel, and hence only one service at a time over each of the channels.

There is no standard for the SPID number; however, many ISDN providers implement the standard seven-digit telephone number format as the SPID number.

ISDN call setups are highly dependent on the type of switches and services provided by the carrier's central office. The United States primarily uses the Signaling System 7 (SS7) switch in the central offices. The SS7 inter-office communication is at 64 Kbps. Prior to the SS7 switch, the inter-office communication was at 56 Kbps. If the ISDN connection is specified to an office without an SS7 switch, the router must be configured to use 56 Kbps when placing a call and have the bandwidth of the ISDN interface manually configured to 56 Kbps.

Another case a designer must be aware of is when calls are made at 56 Kbps, but the receiving end clocks at 64 Kbps. This incompatibility of the line speeds can cause data corruption and the ISDN setup to fail. The Cisco routers can be configured to handle this change in line speeds of a completed call by automatically clocking the line at the line speed received on the incoming call.

Using some of the advanced services provided by the carriers, a router can perform some connection time security using caller ID. The central office switch must support the caller ID feature, and the router must have the screening ISDN security feature enabled. When a call is received, the router interprets the called ID provided in the setup message. The caller ID is matched against configured incoming caller ID values. If the incoming caller ID does not match, the router rejects the call.

PRI and BRI

The ISDN primary rate interface (PRI) and basic rate interface (BRI) provide the physical connection to the ISDN network. Each PRI and BRI bundles B and D channels. The bearer (B) channel is rated at 56 or 64 Kbps for data or voice services. The D channel provides the signaling and controls of the services for ISDN end-to-end connectivity establishment. In some instances, the D channel is used as a low-bandwidth data channel and operates at 16 Kbps.


The BRI ISDN connection provides two 64-Kbps B channels and one 16-Kbps D channel. This is also referred to as a 2B+D line. The two 64 Kbps B channels can establish connections to different locations, one can be used for data and the other for voice, or they can be "bonded" to act as a full 128-Kbps data connection between two points.

The PRI ISDN service is provided over T1- or E1-leased lines from the customer premise equipment (CPE) to the ISDN switch at the carrier central office. The T1-based PRI provides 23 B channels and one D channel. The D channel on a PRI-based T1 is the 24th channel on the circuit. The E1-based PRI service, mostly found outside North America, provides 30 B channels and a 64-Kbps D channel (30B+D). The D channel on a E1-based PRI service is the 16th channel on the E1 connection.

DDR Model

Cisco DDR uses a design stack for providing DDR services that is comprised of five layers (see Figure 11-1). Although DDR is not an actual networking architecture, the use of a design model enables network designers and engineers to build scalable and fault-tolerant DDR internetworks while meeting performance, service, and cost requirements.

Figure 11-1
DDR design model.



DDR Dialer Cloud

Cisco DDR supports switched circuit connections over the following interface types:

- Integrated Services Digital Network (ISDN) interfaces using BRI and PRI
- ISDN terminal adapters (TAs)
- Asynchronous serial interfaces available on the auxiliary port on Cisco routers
- V.25bis and DTR dialing for switched 56 CSU/DSUs
- Synchronous modems

The switched circuit connections use one of the following protocols or encapsulation techniques to establish end-to-end transport:

- PPP
- HDLC
- X.25
- SLIP

The type of dial-interface and protocol used depends on the DDR services required. Establishing the services is based on the dialer cloud created by the DDR internetwork design.

The dialer cloud is the network formed by the interconnection of DDR devices. A dialer cloud is a network concept that is collectively the potential interconnections and the current active point-to-point connections. The active connections have the characteristics of a non-broadcast multi-access (NBMA) media network akin to frame relay networks. The outbound connections used to establish the dialer cloud and the router must map the network protocol address to a directory number indicating the partner(s) of the dial-out process. The inactive DDR connections are "spoofed" by the router, so as to allow them to appear in the routing tables exchanged over the WAN dedicated links.

The dialer cloud is directly affected by the design of the DDR internetwork. Therefore, it is important to understand the different topologies, traffic, protocol addressing in use, routing structure, and the mechanisms used for triggering the DDR connection.

IP Addressing

Since a dialer cloud provides network communication between locations, an addressing scheme must be established for the various protocols that will be mapped to the dialer cloud. The router configuration maps protocols to the dialer interface used as the DDR connection. The two methods used for assigning the networking addresses to the dialer interfaces are subnet and unnumbered.

Subnet allocation to the dialer cloud amounts to a unique address shared by all the DDR routers that enter the DDR dialer cloud. This type of addressing is really no different than that used now in a LAN or multipoint WAN, or an NBMA networking environment. The addressing scheme used by one interface must be followed through on all the other dialer interfaces used in the dialer cloud.

Using unnumbered interfaces is similar to the unnumbered addressing implemented on Cisco WAN interfaces. Unnumbered addressing uses the address of another interface as the address of the dialer interface. This is possible, due to the fact that the routing table will point to the dialer interface and a next-hop address, which is represented by the dialer interface mapping.

Topology

The design of the DDR internetwork is closely related to the traffic requirements and the number of sites supported on the dialer cloud. In determining the type of topology to plan for, a network designer must discover the frequency of traffic required between DDR locations. The traffic analysis applied here will also assist the network designer in determining which DDR site is to initiate the DDR connection. The three basic topologies used for DDR are point-to-point, hub and spoke or star, or a fully meshed topology.

Point-to-Point

The simplest of DDR topologies is the point-to-point configuration, shown in Figure 11-2. Each site has defined a dialer interface that maps the remote sites address to a telephone number. For instances when extra bandwidth can be the impedance of a DDR connection, multiple DDR connections can

be established where the connections use Multilink PPP to aggregate the total bandwidth as a single link between the sites.

Hub and Spoke

A hub and spoke or star topology is built on the notion of DDR connections occurring on a central site that manages every remote DDR connection. This is typically found when DDR is used as a backup topology for routers that have lost their primary dedicated WAN link. The central site or hub of the configuration is often the corporate telecommunications data center. Figure 11-3 illustrates this scenario.

The hub and spoke topology provides for centralized management of the DDR cloud and therefore is easier to configure on the routers at the remote locations. Any-to-any connectivity between the remote locations occurs via traditional routing functions, rather than establishing a new DDR connection to the remote location.

Figure 11-2
DDR point-to-point topology.





Figure 11-3
Hub and spoke topology for DDR internetworks.



Fully Meshed

Fully meshed configurations, as shown in Figure 11-4, lend themselves to providing direct any-to-any connections, yet they increase the complexity and cost of implementation. Each DDR router must have the complete dialer mapping for connecting to all the other remote DDR routers. Providing this capability means careful coordination between switched circuit providers and adequate dialer interfaces to support the potential for all remote DDR routers connecting at the same time to any other DDR router.

Dial Service Considerations

The media used for the DDR connection is accessed using the Cisco IOS dialer interface feature. The Cisco router platforms support ISDN B channels, synchronous serial interfaces, and asynchronous interfaces as dialer interfaces.

Data Encapsulation

The transported data between two DDR routers must be encapsulated in accordance with the media type used for the connections. DDR supports the following encapsulation techniques:

PPP: Point-to-point protocol (PPP) supports multiple protocols and is available on all the various supported DDR media. PPP also negotiates addressing and provides authentication. Since it is an open standard, it is interoperable between vendors.


HDLC: High-level Data Link Control is supported on synchronous serial and ISDN interfaces. HDLC supports the transport of multiple protocols however it does not have an authentication capability.

SLIP: is only supported by IP protocol on asynchronous interfaces. The addresses used must be manually configured. SLIP does not support authentication.

X.25: supports both synchronous serial lines and single ISDN B channels.

Of these different interface types, the most widely used is ISDN and PPP for the transport of data between the two DDR routers.

Figure 11-4
Fully meshed DDR
network topology.



Synchronous Lines

Synchronous serial interfaces support the dialer feature by using V.25bis dialing or DTR dialing. V.25bis uses in-band dialing, which uses the same bandwidth that carries the data between the connecting locations. The V.25bis standard is used with a synchronous modem, ISDN terminal adapters (TAs), and switched 56 DSU/CSUs. The DTR dialing option is used for only connecting out. If the interface is defined for DTR, connect-in attempts will fail. The DTR signal raised on the physical interface causes the attached device to dial a number. Synchronous serial interfaces can use PPP, HDLC, or X.25 datagram encapsulation techniques.

Asynchronous

Asynchronous DDR connections are made through communications servers or by using the auxiliary port on the router. DDR connection setup using asynchronous connections takes longer and requires a chat script to initiate the modem dialing and login commands sent to the remote device. Asynchronous connectivity supports outbound calls and, depending on the design criteria, may provide a cost-effective solution for low-bandwidth-switched connections.

ISDN

The ISDN interfaces on a Cisco router are either ISDN BRI or PRI. Typically, BRI is used at remote locations and PRI is used in situations for consolidating BRI access to the DDR dialer cloud. In the hub and spoke configuration, for example, the central site router employs an ISDN PRI interface to support multiple ISDN BRI connections. The two B channels available on a BRI can be aggregated when using Multilink PPP. This doubles the bandwidth of the configured ISDN BRI connection. ISDN interfaces are automatically configured as dialer in-band interfaces capable of connecting out or connecting in. ISDN with DDR supports PPP, HDLC, X.25, and V.120 encapsulation. PPP is the more widely used encapsulation technique for ISDN interfaces.

Dialer Rotary Groups

Multiple DDR interfaces can be grouped together to form a dialer rotary group. The hub-and-spoke or fully meshed topologies can take advantage of this feature. The physical interfaces inherit their configuration from the corresponding dialer interface definition. DDR automatically groups ISDN B channels of a BRI or PRI into a rotary group. Cisco DDR supports dialer groups of multiple physical interfaces, enabling one physical interface to be busy and a second physical interface defined in the rotary group to place or receive a connection.

Dialer Profiles

Cisco IOS 11.2 introduced the feature of dialer profiles. Using these profiles, the dialer logical definition is decoupled from a physical dialer interface, thereby adding design flexibility. This feature allows multiple interfaces to use the profile simultaneously, enabling the capability to bridge multiple sites over ISDN.

Dialer Maps

Mapping the protocol address to a telephone number is the core of DDR. This function is called dialer map. The network address is mapped to a telephone number on a DDR dialer interface. Setting up DDR between more than two sites necessitates the use of PPP authentication. Using authentication facilities protocol address mapping to user names, thereby authenticating the connection.

When planning for DDR, it is always good practice to create a table listing the location, protocols needed for dial-in, and the phone number used for making a connection. Using such a table enables the designer to understand the actual connectivity requirements and hence the features required in support of DDR at each location. Using the table, a designer can also quickly note if a location is a dial-in only location, therefore not requiring the dial-out feature support. For dial-in only DDR configurations, the telephone (directory) numbers are not a requirement and are best left out to avoid accidental dialing. Dial-in DDR routers use the PPP authentication name mapping for protocol addresses, ensuring outbound packet placement on the correct PPP connection.

Recent enhancements to Cisco IOS allow dynamic dialer maps using IPCP address negotiation for IP protocol and IPXCP address negotiation for IPX protocol on DDR routers configured as dial-in only.

Routing Packets

The type of routing protocol in use and the frequency of the DDR connection may create an unstable routing environment. To avoid this, DDR routes and directory service tables must be maintained, even during idle time. This is accomplished using the static, dynamic, and snapshot routing techniques.

In addition, default routing and remote node spoofing can be combined with the routing techniques to simplify and stabilize the routing design of the DDR dial cloud.

Static Routing

Static routes defined for connecting the DDR dialer cloud routers are entered manually on each DDR router. Using static routes eliminates the need for DDR routers to broadcast routing updates and tables over the DDR connection. Because static routes are not broadcast, we typically redistribute the static routes, defining the path to the DDR router connection into the dynamic routing protocol used for the router-dedicated backbone. Redistributing these routes ensures DDR route awareness for all the routers of the enterprise-wide network.

Dynamic Routing

Dynamic routing protocols are RIP, RIP2, EIGRP, IGRP, and OSPF. When using these protocols, routing updates flow between the DDR routers once the DDR connection is made. Typically, only the distance vector routing protocols of the dynamic routing protocols are used for dynamic routing in a DDR dialer cloud. This is because link-state routing protocols like OSPF end up performing convergence every time a DDR connection is established or broken. To alleviate this problem, OSPF has an on-demand circuit feature to take care of the DDR connection.

Using Passive Interfaces

The Cisco IOS feature called passive interface prohibits the selected dynamic routing protocol from sending route table updates over an interface. Making the DDR dialer interface passive prevents the router from making a DDR connection based on route table updates. Note that if the DDR is made active, routing updates are still prohibited from traversing connection; however, the router receives route table updates on the DDR connection.

The exception to this is EIGRP and OSPF, which do not accept routing updates when an interface is specified as passive. This type of scenario can be useful in a hub-and-spoke topology when the remote locations are set to

passive and the central DDR router advertises routes over the DDR connection.

Split Horizons

For DDR connections that support two or more B channel connections over a single physical interface, the split horizons function should be enabled. Split horizon is a function of the router used with distance vector routing protocols to reduce routing loops. Enabling split horizons on the DDR connection prohibits the router from advertising routes learned on an interface back onto that interface.

Enabling split horizon in a hub-and-spoke configuration will prohibit remote locations from communicating with each other through the hub. In such a case, split horizon should be disabled on the single physical interface. Cisco has subinterface features, however, that can be applied to ISDN PRI lines that enable each BRI channel used on the PRI to be represented as a unique physical interface, allowing split horizon without affecting remote-to-remote communications.

Dynamically Connected Routes

There are two possible methods for having dynamic connected routers: per-user and PPP peer routers. Per-user AAA-installed routers are defined on AAA servers and are associated with users defined in the AAA server. Using per-user AAA provides a means of authenticity and protection. PPP peer routers use IPCP address negotiation by defining the full host address as a subnet mask for the remote peer. This route is then propagated to the rest of the network through normal routing protocol convergence.

Snapshot Routing

In conjunction with dynamic routing, snapshot routing, or using a client-server model, controls the routing table updates between the DDR routers. Shown in Figure 11-5, when using a hub-and-spoke configuration, the hub is defined as the snapshot server and one or more DDR connecting routers are defined as snapshot clients. Snapshot routing works with the following distance vector routing protocols:

- Routing Information Protocol (RIP) for IP
- Interior Gateway Routing Protocol (IGRP) for IP
- Routing Information Protocol (RIP) and Service Advertisement Protocol (SAP) for Novell Internet Packet Exchange (IPX)
- Routing Table Maintenance Protocol (RTMP) for AppleTalk
- Routing Table Protocol (RTP) for Banyan VINES

During DDR connection, the client and server exchange routing tables and then go into a quiet period. At the end of the defined quiet period, the client and server again exchange routing tables. This type of routing table update mechanism preserves valuable bandwidth since it is common that distance vector routing protocols transmit their entire routing table in 10- to 60-second intervals. Snapshot routing therefore protects the DDR bandwidth from being overutilized every 10 to 60 seconds for routing updates.


Snapshot is very useful in networks that have few topology changes and a stable physical topology. In networks that use OSPF or EIGRP as the backbone routing protocol, standard route redistribution can be used to propagate the updates between the routing protocols.

The snapshot client router ISDN interface configuration specifies the key variables for enabling snapshot routing:

- The length of the active period (which must match the length specified on the central router)
- The length of the quiet period, which can be specified up to 100,000 minutes or 69 days
- Whether the router can dial the central router to exchange routing updates in the absence of regular traffic
- Whether connections that are established to exchange user data can be used to exchange routing updates

Employing snapshot routing is a good way for periodic testing of the ISDN backup line. Setting the snapshot parameters and DDR triggers properly enables management to address any DDR line problems before they are needed in a backup situation. Although snapshot is used primarily for keeping remote location DDR router table updates in synch with the backbone network, it can be used during active backup purposes as a means of reducing bandwidth consumption during backup and recovery.

Figure 11-5
Snapshot-routing
client-server model.



DDR as Dial Backup

DDR used as a dial-backup mechanism for dedicated WAN links requires planning for the use of floating static routes or backup interfaces. Typically, the network designer must address the total number of backup lines to support at any given time and plan the appropriate router interface support. Usually, a hub-and-spoke configuration is useful for DDR as a dial backup. In such scenarios, the hub is setup as dial-in only and the remotes are set up as dial-out only. In this way, contention avoidance is achieved during recovery.

Connection Triggering

The DDR connection established by the router can be triggered by various set parameters. In some instances, a load threshold can be set to establish a second DDR connection between two routers for extra bandwidth. In other scenarios, packets can be identified as "interesting," which will be

interpreted to establish a DDR connection to another DDR router for a period of time.

Bandwidth

DDR enables a utilization threshold to be specified on a DDR connection. If this utilization is met, a second DDR connection is established to the remote DDR router, thus meeting performance requirements. This load-threshold parameter is calculated based on the bandwidth value defined for the dialer interface being used. The load is calculated dynamically.

If the specified threshold is met, DDR then establishes a second DDR connection to that remote router. The Cisco router distributes packets over multiple DDR links between two DDR routers based on the queue depth. The link with the shortest queue receives the next packet for sending. If a scenario arises in which the load-threshold value is met for a period of time, the router establishes a second DDR connection to meet the performance requirements.

Using this technique with a low load-threshold specified may result, however, in a single DDR connection being used for transport, since there is no longer a queue depth, but costs are incurred for establishing the second link for the duration of the DDR session. This can occur on BRI connections using the two B channels, as a DDR link is between two routers, and PRI lines use a round-robin technique, ensuring that multiple DDR connections are used. Idle secondary DDR connections are avoided by ensuring that the load-threshold value is not set too low. It is recommended that 50 to 70 percent of the line utilization be set as the trigger for a second DDR connection.

Interesting Packets

Packets are defined as interesting or uninteresting, and interesting packets trigger DDR connections. Uninteresting packets are those that are not destined for active interfaces on the router and will not trigger a DDR connection. The router therefore drops uninteresting packets, while interesting packets are those that have a destination outside the router active interfaces and will trigger a DDR connection. Packets are defined as interesting packets by using a permit access list, while uninteresting packets are defined using a deny access list. The following list outlines the decision process.

1. If a packet is uninteresting and there is no active DDR connection, the packet is dropped.
2. If a packet is uninteresting and there is an active DDR connection to the specified destination of the uninteresting packet, the packet is sent over the DDR connection.
3. If the packet is interesting and there is no active DDR connection to the destination of the interesting packet, a DDR connection will attempt to be made.
4. If an active DDR connection exists for a destination specified within an interesting packet, the packet is sent and an idle timer is then reset.

Controlling Routing Updates as Triggers

A router for any routing protocol supported on the router after waiting a period of time typically sends routing updates. Distance vector routing protocols send their entire routing table periodically. Link-state and advanced routing distance routing protocols (such as Enhanced IGRP) send hello messages to verify the connection to another router using the same protocol. Table 11-4 lists the various default periods for periodic routing update cycles of IP networks.

In the case of OSPF, IS-IS, and Enhanced IGRP, the DDR connection for the destination of the hello messages keeps the DDR connection up all the time. This is not only costly but also a waste of router resources. For distance vector protocols like RIP and IGRP, the DDR connection will “flap” between connected and disconnected, causing excessive route convergence on the partner DDR routers and any routers connected to them. It is for this reason that it is recommended to apply a deny access list to the above routing messages listed in the table.

Table 11-4

Periodic IP Routing Protocol Messages

| IP Routing Protocol | Update Period |
|---------------------|--|
| Enhanced IGRP | Five seconds (Hello) |
| IGRP | 90 seconds |
| RIP | 30 seconds |
| OSPF | 10 seconds (Hello, depends on link type) |
| IS-IS | 10 seconds (Hello) |

Note that once a DDR is active, these routing messages are not filtered and flow through to the destination DDR routers. In this way, it is the application requirements of the location that dictate whether the DDR connection is made and not the inherent network protocol requirements that trigger the DDR connection. Once a DDR connection is made, however, the inherent network protocol mechanisms are used to establish routes to the destinations.

Novell IPX networks characteristically generate many different types of packets on a periodic basis. These too should also be reviewed and marked as uninteresting to avoid the flapping of DDR connections and the costs in router CPU cycles and connection times. The packets that should be reviewed along with their default update time are listed in Table 11-5.

The watchdog packet is used by a server to verify that a previously established client is still on the network and is maintaining a session with the server. This packet is sent out approximately every five minutes. Instead of using an access list filter to deny this packet, the Cisco IOS employs a spoofing technique. Using the watchdog-spoofing technique, a router replies back to the server on behalf of the destination IP workstation. In this way, the DDR connection is not established for this keepalive packet between the server and the client.

Access Lists

Cisco IOS access lists are used in IP network environments to trigger the DDR connection for any type of IP application. The access list used can define a specific IP host or it can be granular in support of a TCP or UDP port number for activating the DDR link. A generic access list can be created that can be used for all the DDR connections, instead of a unique access-list for each interface. In this way, the access list can be applied to any interface, as it is required. Similarly, access lists to deny or permit on Novell packet types can be defined for generic denial and granular permits.

Table 11-5

Novell IPX Packet
Types Requiring
Attention Using
DDR

| Novell Packet Type | Update Period |
|--------------------|---------------|
| RIP | 60 seconds |
| SAP | 60 seconds |
| Serialization | 66 seconds |
| Watchdog | Five minutes |

Security

Security for switched connections is a great concern. Many corporate networks have been compromised due to a lack of authentication by a connecting resource. Cisco has two basic front-line authentication features available for ISDN implementation, callback and screening. This section briefly discusses each of these options.

Callback

Using the callback feature, a router receiving a call from another location is requested to callback the calling router. Typically, a central site router receives a call from a remote location and calls back to the remote location using PPP along with AAA authentication.



Screening

If caller ID is supported by the central office ISDN switch, a router receiving an incoming call can be screened. Using the caller ID found in the setup message, the router can determine if the calling number is a valid participant for making connections to this router. If it is, the call setup continues. If the caller ID presented by the ISDN switch is not found, the receiving router disconnects the call.

CHAPTER

12

Preparing the Cisco Router



This chapter discusses the basics of router installation and the core parameters that reflect loading the IOS code, configuring the router, loading microcode, and managing the files.

Determining the Proper IOS Code

The Cisco IOS system image distribution system is quite extensive and at times confusing. Prior to deciding on what image is required for each router, it is important to understand the services required at each location. The quick checklist needed to answer these questions follows:

- Which network protocols are required to facilitate business processes at each location?
- Which, if any, network protocols are unique for site-to-site location?
- Is temporary connectivity required?
- What is the addressing plan for each network protocol?
- Which routing protocol(s) will be in use to facilitate connectivity?
- What are the WAN protocols needed for each interface?

Once you have answers to these questions, you must match them with the different feature sets available for the hardware platform. The feature sets are categorized into the following:

- *Basic*: The basic feature set for the hardware platform.
- *Plus*: The basic feature set plus additional features, depending on the hardware platform selected.
- *Encryption*: The addition of 40-bit (Plus40) or 56-bit (Plus56) data encryption feature sets to either a basic or plus feature set.

The Cisco IOS releases vary and fluctuate with code fixes almost weekly. Table 12-1 is included to assist you in understanding the position in the release lifecycle of a Cisco IOS release.

Two types of images can be found on Cisco routers. The system image is the complete Cisco IOS software that loads into the router RAM and is used to operate the router. The second type of image is the boot image. The boot image is a subset of the Cisco IOS software and is used to load the complete Cisco IOS software image on the router at start-up or on execution of the Cisco enable configuration command `BOOT`. For all Cisco router platforms except those specified in Table 12-2, the images are located in Flash memory.

Table 12-1
Table Listing Cisco
IOS Release Generations

| Release Type | Description | Timing | Numbering Example |
|---|---|---|--|
| Major release—Functionally Complete Software (FCS) | Introduces significant features, functionality, and/or platforms on a stability-oriented release vehicle | As needed to support customer needs | 12.0(1) |
| Major release—Scheduled maintenance updates | Periodic revisions to major releases: fully regression-tested, incorporates the most recent bug fixes, and no new platforms or features—focused on stability | Regular maintenance cycles | 12.0(3) |
| Major release—Interim builds | Working builds, usually not regression-tested, and not intended for customer use except in unusual circumstances | Weekly | 12.0(4.2) |
| General deployment (GD) | A major release that is appropriate for general, unconstrained use in customers' networks | When stability of release has been proven internally by Cisco and externally by customers | 12.0(8) and all subsequent maintenance updates of 12.0 [12.0(9), 12.0(10), etc.] |
| Early deployment (ED)—FCS | Introduces significant new features, functionality, and/or platforms on a feature-oriented release vehicle: based on a major release and will not achieve general deployment | As needed to provide support for newly emerging technologies | 12.0(1)T |
| Early deployment—Scheduled maintenance updates | Periodic revisions to ED releases: fully regression-tested, incorporates the most recent bug fixes including those from Major Release, and usually delivers new platforms and/or features | Regular maintenance cycles | 12.0(3)T |
| Early deployment—Interim builds | Working builds—usually not regression-tested and not intended for customer use except in unusual circumstances | Generally weekly, though some ED Releases may follow a different policy | 12.0(4.2)T |

Table 12-2

System and Boot Image File Locations for the Cisco 7000, 3600, and 1600 Series Routers

| Router | Flash (flash:) | Bootflash (bootflash:) | First PCMCIA Slot (slot0:) | Second PCMCIA Slot (slot1:) |
|-------------------|----------------|------------------------|----------------------------|-----------------------------|
| Cisco 7000 family | - | yes | yes | yes |
| Cisco 3600 series | yes | - | yes | yes |
| Cisco 1600 series | yes | - | - | - |

The image name format identifies the router platform, features, and the type of area on the router from which the image executes. The image name is formatted as *platform-features-type*.

Table 12-3 provides a sample listing of the platform variable of the IOS naming convention.

The feature variable identifies the feature sets included in the image. Table 12-4 lists some examples of the feature variable used in the IOS naming convention. If more than one feature set is included, the features are listed in the variable by alphabetical sequence.

The *type* field identifies the location of the running image in the router.

Table 12-3

Sample Platform Variables for IOS Image Name

| Cisco Router Platform | platform Value in IOS Name |
|--------------------------------|----------------------------|
| Cisco 7000 series with RSP7000 | rsp |
| Cisco 7500 series | rsp |
| Cisco 4500/4700 series | c4500 |
| Cisco 4000 series | c4000 |
| Cisco 3600 series | c3600 |
| Cisco 2600 series | c2600 |
| Cisco 2500 series | c2500 |
| Cisco 1600 series | c1600 |
| Cisco 1005 series | c1005 |

Table 12-4
Sample Feature
Variable Values for
IOS Naming Con-
vention

| Variable | feature Value in IOS name |
|---|------------------------------|
| APPN | a |
| ATM | a2 |
| Desktop subset (SNMP, IP, Bridging, WAN, Remote Node, Terminal Services, IPX, Atalk, ARAP) (11.2—Decnet) | d |
| Reduced desktop subset (SNMP, IP, IPX, ATALK, ARAP) | d2 |
| IPeXchange (no longer used in 11.3 and later) StarPipes DB2 Access enables Cisco IOS to act as a “gateway” to all IBM DB2 products for downstream clients/servers in 11.3T | e |
| FRAD subset (SNMP, FR, PPP, SDLLC, STUN) | f |
| ISDN subset (SNMP, IP, Bridging, ISDN, PPP, IPX, Atalk) | g |
| Enterprise subset (formerly bpx, includes protocol translation), <i>not used until 10.3</i> | j |
| kitchen sink (enterprise for high-end) (same as bx) (Not used after 10.3) | k |

This *type* field can be one of the following values:

- *f*: The image runs from Flash memory.
- *m*: The image runs from RAM.
- *r*: The image runs from ROM.
- *l*: The image is relocatable.
- *z*: The image is zip-compressed.
- *x*: The image is mzip-compressed.

Locate IOS Using Cisco Connection Online (CCO)

The Cisco IOS software image selected for the router can be accessed directly from the Cisco Systems Web site named Cisco Connection Online (CCO). The CCO Web page is at <http://www.cisco.com> and the Cisco

IOS software images are found at <http://www.cisco.com/kobayashi/sw-center/sw-ios.shtml>. Here the supported and available Cisco IOS release levels are listed. Drilling further down by selecting a release level, the specific router model and feature sets are selected.

CCO Software Center

Authorized users with a user ID and password access the CCO software center. Typically, these are provided after you have purchased a maintenance agreement with Cisco Systems. To locate the release, you must follow the following path through the software center. In this example, we are locating a major release of Cisco IOS 12.0:

```
http://www.cisco.com/
```

Enter the following in the browser:

```
http://www.cisco.com/kobayashi/library/12.0/index.shtml
http://www.cisco.com/cgi-bin/iosplanner/iosplanner.cgi?major-
Rel=12.0
LOGIN to CCO
```

From here, select the hardware platform of the Cisco router being prepared for upgrading. After selecting the router platform, select the feature set required, based on your answers from the previous questions. At this point, CCO Software Center will require you to verify the agreement for downloading the software and then display the name of the Cisco IOS software image to be downloaded to your computer. Upon selecting the software image, a software license agreement appears to which you must answer yes to obtain the IOS software image. The image name is then shown with three options for downloading to your location:

- FTP directly to your computer
- Use HTTP to download the image
- Receive an e-mail message with an attached file

If these approaches do not work for any reason, the same IOS software image can be obtained by executing a FTP connection to <http://www.cisco.com> and logging in with the registered CCO user ID and password. In this case, the image name found from the previous CCO software center steps outlined above are entered using the FTP “get” command. For the example we used, the FTP command will be entered as:

```
get /cisco/ios/12.0/12.0.1a/7500/rsp-jsv-mz.120-1a.bin
```

In either case, be sure to specify a directory on your computer that has enough storage space to handle the IOS software image.

Downloading IOS to the TFTP Server

Typically, the IOS software image retrieved from CCO is stored on the UNIX computer or a Windows 95/98/NT workstation that executed the download. Cisco routers, in a corporate environment, are typically updated using Trivial File Transfer Protocol (TFTP). Using any number of free TFTP server offerings on the Internet for a Windows 95/98/NT workstation, you can download the IOS software image to the target router.

The following steps must be followed with either a UNIX or Windows platform:

1. Start the TFTP server application on the UNIX computer or Windows workstation.
2. Ensure the correct IOS image is placed in the TFTP directory.
3. Be sure to use the exact IOS image name, as found on the TFTP server.


The last step is important because some browsers and/or operating systems may change the name of a .bin extension to a .exe extension on the file name. Once these steps have been verified, the IOS can be transferred to the target router.

Loading IOS on the Router

The IOS image stored on a TFTP server must be written to the target router. Although the router supports alternative methods for transferring a system image, TFTP is the most widely used. FTP transfer of the IOS image is available, starting with Release 12.0 of Cisco IOS. A Cisco router arrives with an onboard Cisco IOS system image for initial start-up. Typically, routers are ordered with flash memory to allow the storage of new IOS software images. Access to the contents of the flash memory requires the end user to be in privileged mode.

Figure 12-1

File transfer of Cisco IOS image to and from a TFTP or FTP server.



NOTE: Cisco routers do not support remote file transfer initiation. All transfers of files on a router must be initiated from the router.

Enter Privileged Mode on the Terminal Interface

Privileged mode on a Cisco router allows the end user to access system-level commands that can affect router or network performance as well as configuration mode commands. Entering privileged mode is accomplished by entering the EXEC mode command:

```
enable
```

The router prompts for a privileged password at this point. If this router is starting up for the first time, just entering the enable EXEC mode command will allow the end user to enter the privileged mode.

Once the end user has entered privileged mode, the IOS copy command can be executed to either get an IOS image from a server or to put an IOS image on a file server.

Issue Copy to/from a TFTP Server

Before loading a new IOS on an existing Cisco router, it is prudent to write the previous IOS image used on the router to a server. Existing routers typically have the onboard system image that came with the router and a system image in the flash memory of the router.

Copying an IOS image from a 7000, 7200, or 7500 series router to a TFTP server requires the use of the following command:

```
copy device:filename tftp:[[//location/]directory/]filename]
```

The *device:filename* variable is the value of the flash memory location on the 7000, 7200, or 7500 series router. An example of a resolved *device:filename* variable is `slot1:ios12.0-1`, which indicates that the file named `ios12.0-1` on the flash PCMCIA memory card, located in slot1 of the primary RSP card on a 7500 series router, is the file to copy. The location, directory, and filename positional optional variables of the **tftp** keyword specify the TFTP host name or IP address, the directory tree to use on the TFTP server, and the name of the file to be written on the TFTP server.

For example, copying an image from a 7000 router to a TFTP server using the image found on the internal flash of the 7000 router is entered as

```
copy flash:ios12.0-1 tftp://tftpserver/ciscoimages/7000/ios12.0-1
```


Optionally, this command can be abbreviated and the router can prompt you for the variable inputs. For instance, the above command can be entered as

```
copy flash: tftp:
```

and the router will prompt the end user with the following questions:

```
IP address of remote host [255.255.255.255]?
filename to write on TFTP host?
```

The router administrator must provide the IP address of the TFTP server or the DNS name of the TFTP server, and then the IOS image name found (in this case) in the flash memory location. In this instance, the router writes the IOS image specified to the TFTP server default directory.



NOTE: When writing files to a server, the process must have the capability to create a new file. If not, the copy will fail. To avoid this, either set the processes with the authority to create files or have a predefined file with the same name already on the server.

The value for the *device* parameter of the *device:filename* variable is listed in Table 12-5 with a description for each variable.

On routers that support multiple flash memory locations, the change directory (`cd`) command must be entered, pointing to the flash location. This is similar to the change directory command for any UNIX or Windows command. For example, specifying `cd slot0:` changes the working directory from flash to the PCMCIA memory card installed in memory slot0 of the router. In the case of a Cisco router with multiple flash memory locations, the command points the copy command to the correct location just by specifying copy.

Some routers, like the 3600 series, allow the partition of flash memory. In this case, the copy command must be entered by specifying *device:partition:filename* to point to the appropriate flash memory location. For example, writing an image from a 3600 series router to a TFTP server can be accomplished using the following **copy** command:

```
copy slot0:2:c3600-ios tftp:
```

In this example, the flash memory card in slot0 is accessed and the file `c3600-ios` is found in partition 2 of the flash memory card installed in slot0 of the router. The router will prompt for the destination TFTP server and then verify the write command.

Table 12-5

The Device Parameter Values Available for the **device:filename** Variable for Copying IOS Images to a TFTP Server

| <i>device</i> Parameter Value | Description |
|-------------------------------|--|
| flash | Internal flash memory on the router |
| bootflash | Internal flash memory specific to the 7200/7500 series routers |
| slot0 | PCMCIA flash memory card in the first or only slot on the router |
| slot1 | PCMCIA flash memory card in the second slot on the router |
| nvrn | Internal non-volatile RAM on the router |
| slavebootflash | Internal flash on the 7507/7513 slave RSP card configured for high availability |
| slaveslot0 | First PCMCIA flash memory card on the 7507/7513 slave RSP card configured for high availability |
| slaveslot1 | Second PCMCIA flash memory card on the 7507/7513 slave RSP card configured for high availability |
| slavenvrn | Internal non-volatile RAM on the slave RSP card of a 7507/7513 configured for high availability. |

Issue Copy to/from the FTP Server

The FTP for transferring files between the router and a server was introduced with Release 12.0 of Cisco IOS. Since FTP is a connection-oriented application using TCP/IP, it provides better throughput and a higher rate of success than does TFTP, which is a connectionless oriented applications using UDP/IP for delivery services.

Being a connection-oriented application, FTP requires the use of a login and password prior to transferring the IOS image. Establishing the username for the login and the password presented by the router to the FTP server is accomplished using the configuration operation of the privileged mode on the router. The following sequence is used to again provide this information to the router IOS in preparation of using FTP to transfer an IOS image.

```
enable
Enter password: xxxxxx
configure terminal
ip ftp username login-name
ip ftp password login-password
end
write memory
```

As with TFTP, the **enable** command, followed by the privileged password, allows the operator to gain access to privileged mode commands. The **configure terminal** command executes the configuration mode of the router. In this mode, any and all options and definitions can be entered to modify, add, or delete the router running configuration. The FTP username associated with a login name on the target FTP server is provided to the router IOS by the **ip ftp username** command. The **ip ftp username** login-name variable defines the default user login name used by the router if a login name is not specified on the **copy ftp:** command.

The *login-name* variable is a valid user name on the target FTP server. The **ip ftp password** command is the default password associated with the specified *login-name* on the target FTP server. The **end** command exits the configuration mode.

After executing these commands, the router IOS configuration has the FTP parameters necessary to connect to a FTP server. The **write memory** command at the end of the statements saves the running configuration in the router to the router's NVRAM for permanent storage. Saving the running configuration to memory saves the just entered configuration information between power-on reset and router reloads.

Backing up the current IOS image to an FTP server is accomplished by entering the following command:

```
copy device:filename ftp:[[//[/login-name[:login-password]@]/location]/directory]/filename]
```

Using this command, the router administrator specifies the flash location of the router using the *device* parameter of the *device:filename* variable and the name of the image using the *filename* parameter. The *device* parameter can be any value specified in Table 12-5 and the *filename* is an image name found on the flash device.

The *login-name* and the *login-password* values may be the values set by the **ip ftp username** and the **ip ftp password** commands. However, remember that these commands are used only for default. If the commands were never entered, the FTP server prompts the router administrator for the user name and password. The remaining variables of location, directory, and filename are defined the same as with TFTP copy.

For example, to transfer an image from flash memory to an FTP server at the server name FTPSERVER, the router administrator enters the following command:

```
copy flash:ios-image ftp://joev:jandj@FTPSERVER/cisco/image/ios-image
```

The image named ios-image is copied to the directory //cisco/image/ in relation to the directory structure assigned to the FTP server authorized user named joev. The IP address of the server named FTPSERVER is determined by the router performing a DNS query to the specified DNS server found in the router configuration file.

Transferring a new image file to the router is performed using the following copy ftp: command:

```
copy ftp:[[//[/login-name[:login-password]@]/location]/directory]/filename] device:[filename]
```

The **copy ftp:** command uses all the same variables as previously discussed. In using the command, you need only specify the following:

```
copy ftp: slot1:
```


Using this abbreviated format of the command, the default *login-name* and default *login-password* are passed to the FTP server. If the router does not have the defaults specified, the FTP server prompts for the login name and the login password. The FTP connection is treated like any other and the router administrator then enters the *filename* of the file to retrieve. Note that, in this instance, the default directory structure associated with the login name must have the requested *filename* to retrieve. The requested file will be saved on the flash memory card specified as slot1 in this example.

The Cisco IOS uses a default mechanism for providing the FTP *login-name* and *login-password* values. For the *login-name*, the IOS will use the following criteria:

1. The *login-name* specified with the **copy** command, if a *login-name* is specified.
3. The *login-name* set by the **ip ftp username** command, if the command is configured.
3. The default FTP login name of anonymous.

The password is determined by the following:

1. The *login-password* specified in the **copy** command, if a *login-password* is specified.
2. The *login-password* set by the **ip ftp password** command, if the command is configured.
3. The router forms a *login-password login-name@routername.domain*. The variable *login-name* is the login name associated with the current session, *routername* is the configured host name of the router, and *domain* is the domain name of the router.



NOTE: Copy operations can be interrupted by pressing **Ctrl-^** or **Ctrl-Shift-6**. This terminates the current copy operation, but the partially copied file remains in flash memory until erased.

Loading CIP or CPA Microcode on a Cisco 7000/7200/7500 Router

The process of transferring CIP/CPA microcode to the 7000/7200/7500 (7K) series routers uses the same mechanisms as that of transferring IOS system images. Either `copy tftp:` or `copy ftp:` commands can be entered to retrieve the CIP/CPA microcode. The CIP/CPA microcode images once copying begins are self-exploding images. During a copy operation, the following output occurs on the terminal:

```
routername#copy tftp slot0:
Enter source file name: cip22-25.exe
6843800 bytes available on device slot0, proceed? [confirm]
```

```

Address or name of remote host [tftp.domain.com]? 10.16.47.72
Accessing file "cip22-25.exe" on 10.16.47.72 ...FOUND
Loading cip22-25.exe from 10.16.47.72 (via Ethernet2/2): !
-- expanding multi-segment file --
slot0:cip22-25.exe_kernel_hw4 size = 257888
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!ccccccc
-- expanding multi-segment file --
slot0:cip22-25.exe_kernel_hw5 size = 256914
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!ccccccc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_802 size = 233792
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!ccccccc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_csna size = 85896
!!!!!!!!!!!!!!!!!!!!cc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_eca size = 461408
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!cccccccccccccc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_offload size = 64656
!!!!!!!!!!!!!!!!!!!!c
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_pca size = 69360
!!!!!!!!!!!!!!!!!!!!cc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_push size = 13752
!!!
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_tcpip size = 182032
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!cccccc
-- expanding multi-segment file --
slot0:cip22-25.exe_seg_tn3270 size = 542392
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!cccccccccccccccc

```

In this output, the CIP microcode named cip22-25.exe is transferred from a TFTP server named tftp.domain.com to the router using the **copy tftp:** command. The command specifies that the PCMCIA flash memory card installed in slot0 of the router is the destination for the CIP microcode. As you can see from the example, the CIP microcode is expanded during the transfer process.

Locating CIP/CPA Microcode on the CCO Software Center

In a process similar to locating the IOS image for a Cisco router, the CIP/CPA microcode image is determined. A similar path is followed, as outlined below, to locate the CIP microcode for a 7500 series router using Release 12.0(1) Cisco IOS:

```

http://www.cisco.com
LOGIN to CCO

```



```
http://www.cisco.com/public/sw-center
http://www.cisco.com/public/sw-center/sw-ios.shtml
http://www.cisco.com/kobayashi/sw-center/interworks/cip
```

The location on CCO for the CIP/CPA microcode from the Cisco Web site is specific to the IOS release level. For example, the microcode for Release 11.3 of Cisco IOS is located at

```
http://www.cisco.com/cgi-bin/tablebuild.pl/cip113
```

Release 12.0 of Cisco IOS CIP microcode is found at

```
http://www.cisco.com/cgi-bin/tablebuild.pl/cip120
```

At these URLs, the list of available microcode releases and the supported platform can be found. From here, select the appropriate CIP/CPA code for the router platform being prepared for upgrading. The CIP code is for the Cisco 7000 and 7500 series routers, while the CPA code is for the Cisco 7200 series routers.

After selecting the router platform, select the feature set required, based on your answers from the previous questions. At this point, the CCO Software Center will require that you verify the agreement for downloading the software and then display the name of the Cisco IOS software image to be downloaded to your computer.

After selecting the image, a software license agreement appears to which you must answer yes to obtain the IOS software image. The image name is then shown with three options for downloading to your location:

- FTP directly to your computer
- Use HTTP to download the image
- Receive an e-mail message with an attached file

If these approaches do not work for any reason, the same CIP/CPA image can be obtained by executing a FTP connection to `http://www.cisco.com` and logging in with the registered CCO user ID and password. In this case, the image name found from the previous CCO software center steps outlined above are entered using the FTP “get” command. For example, the FTP command is entered as

```
get /cisco/core/cip/12.0/cip26-4.bin
```

In either case, be sure to specify a directory on your computer that has enough storage space to handle the IOS software image.

Enter config Parameters to Load Microcode

The router configuration file must have a pointer to the CIP microcode name and location. This is accomplished by entering the configuration mode on the router and entering the microcode configuration command:

microcode *interface*[*device:filename* [*slot*] | **system** [*slot*] | **rom** [*slot*]]

The *interface* variable of the microcode command identifies the interface processor board type being addressed in the **microcode** command. Again, the *device* parameter is the flash memory location where the microcode identified by the *filename* parameter is found. The *slot* parameter identifies the slot location of the interface processor type being addressed by the **microcode** command if more than one board of the same interface type is present. Using the *slot* parameter allows you to load different microcode releases into different processors. For example, to specify the loading of CIP microcode CIP22-25.exe, found on flash memory slot0 on the CIP board located at slot 2 on the router, the following command is issued under configuration mode:

```
microcode CIP flash slot0:cip22-25.exe 2
```

This command sets the router configuration to load the CIP card installed on the router with the microcode cip22-25.exe.

After entering this **microcode** command, exit the configuration mode by entering the **end** command. Follow this with the **write memory** or **copy running-config start-config** command to save the new microcode definition in the startup configuration file.

The CIP and CPA interface processors are the only offerings with microcode outside the IOS image. All other processors obtain their microcode from the IOS image used at load time on the router. Table 12-6 lists the possible values for the interface variable on the **microcode** command.

Microcode for interface processors other than the CIP/CPA cards obtain their microcode from the IOS image currently running in the system. Typically, Cisco routers run using the IOS image found in flash memory. The microcode for all processor boards other than CIP/CPA are loaded from this code. The system keyword on the microcode command indicates that the specified interface type should be loaded from the system IOS microcode level.

Loading the microcode is accomplished by issuing the following command under configuration mode:

microcode reload

This command immediately has the IOS reload all the interface processors installed that have microcode running. The **microcode reload** command reloads the interface processors based on the microcode commands entered in the running configuration.

The microcode reload command need only be executed to refresh interface processors or to initially prime an interface processor for the first time. Normal microcode load processing occurs in the following manner:

Table 12-6

The **Interface** Variable Values Available for the **Microcode** Command

| Microcode <i>interface</i> Variable | Interface Type |
|-------------------------------------|---------------------------------------|
| AIP | ATM Interface Processor |
| CIP | Channel Interface Processor |
| EIP | Ethernet Interface Processor |
| FEIP | Fast Ethernet Interface Processor |
| FIP | FDDI Interface Processor |
| FSIP | Fast Serial Interface Processor |
| HIP | HSSI Interface Processor |
| MIP | Multi-Channel Interface Processor |
| POSIP | Packet over Sonet Interface Processor |
| SIP | Serial Interface Processor |
| SP | Switch Processor |
| SSP | Silicon Switch Processor |
| TRIP | Token Ring Interface Processor |
| VIP | Versatile Interface Processor |
| VIP2 | Versatile Interface Processor 2 |

1. The system is booted.
2. A card is inserted or removed.
3. The configuration command **microcode reload** is issued.

Loading the microcode on Cisco 12000 GSR routers is the same as loading images on the Cisco 7500 series routers. In addition, however, each line card on the 12000 series router has a complete copy of the Cisco IOS image. Each line card gets the IOS image after the image is loaded into the GRP of the 12000 series router. The GRP then automatically distributes the image to each line card.

Instances may occur that necessitate upgrading a line card without upgrading the GRP due to temporary fixes that affect the specific function of the line card. To do this, follow the procedures already outlined for the using TFTP or FTP to download an IOS image to a Cisco router. After the download has taken place, the following command is used to load the new image on a specific GSR line card:

microcode {**oc12-atm** | **oc12-pos** | **oc3-pos-4**} **flash** *file-id slot-number*

The first keywords identify the interface type being addressed by the **microcode** command. The *file-id* variable following the flash keyword identifies the IOS image name being used for the load. The *slot-number* variable is the position of the line card on the GSR. By omitting the *slot-number* variable, the IOS image is loaded into all line cards.

The line card is reloaded with the new IOS image by specifying the following command under configuration mode:

microcode reload *slot-number*

The **slot-number** variable is the position of the line card affected by this reload command. If the *slot-number* variable is omitted, all the line cards are reloaded.

Router Basic Configuration and IOS Commands

A few basic IOS commands should be specified within a router for minimal operation. This section discusses these commands and the most commonly used parameters of each command.

Setting up EXEC and Privileged Mode Access

During a first time installation of a router, the router administrator can connect to the router using the console port or the auxiliary port with a personal computer, laptop, or notebook PC. During a first-time installation, the router does not have a password for the privileged mode of the router. In this case, the router administrator simply enters the EXEC mode command **enable** to access privileged mode.

In a previously configured router with a privileged password defined, the router administrator must enter the privileged-password for gaining access to privileged mode. The authorization of the mode allows three attempts. After the failed third attempt, the router administrator falls back to EXEC mode.

Once gaining access to privileged mode, router configuration parameters can be altered. To alter the router configuration, the router administrator must gain access to the configuration mode. Entering the **configuration terminal** command under privileged mode does this.

In first-time installations, it is best to immediately address basic security issues for remote users gaining access to the router. The first of these defenses is the use of a password at the EXEC mode of the router. The EXEC mode is the operational mode presented to all Telnet users and direct serial attached users. Defining this access is performed by using the following basic line commands:

```
line con 0
  exec-timeout 15 0
  password RTRPSWD
  login
line aux 0
  transport input all
  exec-timeout 15 0
  password RTRPSWD
  login
line vty 0 4
  exec-timeout 15 0
  password RTRPSWD
  login
```

Using these configuration commands, three means of access are defined through the console and auxiliary ports and the virtual terminals (vty). The **line con 0** command identifies that the router console port is available for direct access to the router. The **exec-timeout** command following this specifies that the connection will timeout within 15 minutes if data is not transmitted between the router and the end user workstation. The password for EXEC mode access is RTRPSWD and login password authorization is required, meaning that the router administrator gains EXEC mode access using the console port by providing the password.


The **line aux 0** command defines the use of the auxiliary port for access to the router. Typically, this port is used for remote dial-up access to the router in case the router WAN links or LAN connections are down. The auxiliary port in the example uses the same time restrictions, password, and login parameters to authorize remote access to the router. The added parameter **transport input all** enables all possible protocols supported by the auxiliary to communicate with the router. These protocols are shown in Table 12-7.

Table 12-7

The Supported Protocols for the Auxiliary Line on a Cisco Router

| Supported Protocol Value | Auxiliary Port-Supported Protocols |
|--------------------------|------------------------------------|
| lat | DEC LAT protocol |
| mop | DEC MOP remote console protocol |
| nasi | NASI protocol |
| none | No protocols |
| pad | X.3 PAD |
| rlogin | Unix rlogin protocol |
| telnet | TCP/IP Telnet protocol |
| v120 | Async over ISDN |

The final line command is the **line vty 0 4**, which defines five virtual terminal connections. Each terminal connection is used when Telnet connections are established to the router via the LAN/WAN network. All five connections use the same EXEC mode password RTRPSWD to gain access to the basic router operational commands. Each connection will timeout after 15 minutes of idleness.



NOTE: *The **line console** and **line auxiliary** commands allow for only one connection, hence the 0 value following the **con** and **aux** keywords.*

The next time a router administrator accesses the router through any of these connections, the router will prompt for the EXEC mode password. The router administrator has three attempts to enter the newly assigned password RTRPSWD and the router disconnects the Telnet connection after the failed third attempt.

The assignment of a password to the privileged mode is possible using two methods. The first is the use of the **enable password** command. This can be used for providing a hidden password during the printing of the configuration. The **enable password** command is the first stage of a second tier of authorization on the router for gaining privileged access. However, the execution of the **show config**, **show running-config**, or **show startup-config** displays the privileged password on the terminal screen.

Anyone looking over the router administrator's shoulder could easily get the privileged password.

Encrypting the password for display or print requires the inclusion of the **service password-encryption** global configuration command as well as the **enable secret** global configuration command. Using this command in the following format will result in a displayed encrypted command:

```
service password-encryption
enable secret RTRPSWD
```

Using the **enable secret** command, the router encrypts the password RTRPSWD to protect its anonymity. The lower case "s" is used here to demonstrate that the Cisco IOS uses upper and lower case, as well as special characters such as ^, ~, or ` as viable characters for the password. The resulting display from a **show running-config** or **show config** command is

```
enable secret 5 $1$7bcf$LTYfEFUTGj.nrhsG76RBe0
```

The 5 in the resulting display indicates that the encrypted password is shown. In this manner, the displayed configuration on the terminal screen and in print is encrypted.

BOOT SYSTEM List

While in configuration mode, the best practice for insuring the successful loading of a router is through the use of a fault-tolerant booting strategy. A fault-tolerant booting strategy provides a sequential list for the IOS boot strap to locate and successfully load a working IOS image. Using this method, multiple locations can be searched for loading a viable IOS image into the router. The fault-tolerant method can be outlined as follows:

1. Boot the system from flash.
2. Boot the system from a network server.
3. Boot the system from ROM (BOOTFLASH on Cisco 7500).

For example, a fault-tolerant list on a router that has on-board flash and two PCMCIA flash memory cards, along with support for using an FTP server to get the IOS image, would look like the following:

```
Router# configure terminal
Router(config)# boot system flash rsp-IOS
Router(config)# boot system flash slot0:rsp-IOS
Router(config)# boot system flash slot1:rsp-IOS
Router(config)# boot system FTP rsp-IOS 10.6.1.11
```

```
Router(config)# boot system rom
Router(config)# config-register 0x010F
Router(config)# end
Router# copy running-config startup-config
Router# reload
```

Using the above fault-tolerant method, the router uses the list provided based on the boot register value of 0x010F, specified by the **config-register** command. This value indicates that the system will find the IOS image for loading based on the definitions found in the configuration file. The **boot system FTP** command is dependent on the requirements of the FTP server used and whether the **ip ftp username** and **ip ftp password** commands were previously defined in the configuration file. A TFTP server could also have been specified that would not require the login name and an associated password. Because the boot system ftp and boot system tftp commands do not provide a fully qualified file name structure for locating the IOS image, the image must reside in the default directory assigned to the username profile on the server.

Prior to loading the router using the reload command or by executing a power-on reset, the configuration changes made must be copied to the NVRAM of the router in order for the config-register value to take effect. Reloading without performing this step will result in the router performing the locate of the IOS image based on the previous config-register setting.



NOTE: This text assumes that at the end of all configuration changes a write memory or copy running the startup-config command is executed to save the configuration changes in non-volatile memory.

BOOT CONFIG List

Many Cisco router shops have standardized the configuration files such that many of the standard base parameters are the same on all routers. This type of configuration file is known as the network configuration.

Typically, in a network configuration, you will find standardized definitions for routing, bridging, DNS, SNMP, terminal lines, login banners, logging, and possibly access lists. The network configuration file can be merged with a host-specific configuration file. The host configuration file has information specific to a host (router) that is information unique to the individual router. For instance, unique information pertains to interface definitions, access-lists, and host names.

The provision of a configuration list is one that allows for the centralized management of network-wide and router-specific configuration files. For example, the following commands at startup or reload provide a router with the locations for the configuration files specific for the router:

```
Router# configure terminal
Router(config)# boot host tftp://10.16.1.1/tftpboot/hostfile1
Router(config)# boot network tftp://10.16.1.1/tftpboot/networkfile1
Router(config)# service config
Router(config)# end
Router# copy running-config startup-config
```

In the example, the `boot host tftp` and `boot network tftp` commands attempt to retrieve the configuration file from the TFTP server at IP address 10.16.1.1 using the fully qualified file name of `/tftpboot/hostfile1` and `/tftpboot/networkfile1` for the respective commands. If the first attempt does not succeed, the router continues to try every 10 seconds.

The `service-config` command enables the router to use the `boot host` and `boot network` commands to locate the configuration file. If the `service config` command is omitted, the configuration file found in the NVRAM `startup-config` location is used. If the `startup-config` is corrupted and the `boot host` and `boot network` commands are present, however, the `service config` command is automatically enabled to allow the retrieval of a working configuration.

The final step in the procedure saves the running configuration to the NVRAM `startup-config` location. At this point, the router can be reloaded using the `reload` command or a power-on reset.

The format of the `boot host` and `boot tftp` are listed below. The variable parameters follow the guidelines discussed earlier in this chapter. Both of the `boot host` and `boot network` are not necessarily required. Either can be used to load a working configuration for this router at startup.

```
boot host | network {ftp:[[[/[/login-name[:login-password]@]
location]/directory]/filename] | tftp:[[[/location]/directory]/filename]}
```

Again, the FTP login name and login password must match a user profile on the FTP server, if specified.

Managing configuration files is akin to managing the IOS image files. The copy commands previously discussed can assist in moving files from the router to a server or from a server to the router. The only change is the device location of the running and startup configuration files.

In the following example, we protect the current startup config from being changed by first copying it to a flash card on the router. Then we issue a copy of the `running-config` to the `startup-config` location in NVRAM:

```
copy startup-config slot0:old-config
copy running-config startup-config
```

Using this method, we can always manually reboot the configuration pointing at the saved old-config on the flash memory card at slot0, should the new configuration fail. We can do this by issuing a copy of the saved old-config into the running-config location. This will be acted upon as if the commands were being entered on the router in configuration mode.

```
copy slot0:old-config system:running-config
```

Doing this command will restore the router to the previous running configuration.

Assigning a Name to the Router

The router name is very important. Not only does it give the router an entity, its name is used in providing default configuration file names and login names. The router name is defined using the following command:

```
host name
```

The *name* parameter is the name of the router. It is referred to as the host name. It can be used as the entry in the Domain Name System (DNS) server. As we will see later on in the book, a router can have a virtual interface defined that is always considered active. The virtual interface is defined using the **interface loopback 0** command and is accessible from any available LAN/WAN connection that is active on the router. By assigning an IP address to the virtual interface and registering the fully-qualified host name in a DNS with the virtual interface IP address, the router is easily accessible, using the host name for a Telnet connection.

A descriptive name is preferred, one that can provide you with some type of knowledge as to the location of the router. For example, examine the following command:

```
host c1r2Lndn
```

The router name of c1r2Lndn can denote that this router is in the second row in cabinet 1 of the London office. Using a naming convention leads to a simpler troubleshooting solution.

Enabling a DNS Search and Assigning a Domain

The router has Telnet capabilities for accessing other router- or other Telnet-capable hosts. Many times the router administrator may know the name of the host but not the IP address. To facilitate this type of function, the Cisco IOS allows the entry of a domain name and a pointer to the DNS

server used by the network for resolving names to IP addresses. The commands are

```
ip domain-name name  
ip name-server ip-address
```

The *name* variable of the **ip domain-name** command is the fully qualified domain name assigned to the network. The *ip-address* variable of the **ip name-server** command is the IP address of a DNS server for resolving host names to IP addresses. Multiple **ip name-server** commands can be entered to provide multiple DNS servers for resolving the host name to an IP address. Here is a sample of these two commands at work:

```
host c1r2Lndn  
ip domain-name networkcorp.com  
ip name-server 10.10.10.10  
ip name-server 10.10.10.11
```

In this example, the router's fully qualified name is c1r2Lndn.networkcorp.com and the DNS servers that the router will use for resolving host names are found at 10.10.10.10 and 10.10.10.11 IP-addressed host computers. The importance of the **ip name-server** command is not only for Telnet connectivity, but for resolving host names entered on any IOS operational command that may use an IP address. Instead of the IP address, a host name can be used that is then resolved using the DNS servers specified with the **ip name-server** command.

Specifying SNMP for Router Management

Router management is extremely important, especially in large router-based networks. The primary means for managing Cisco routers is through the use of Simple Network Management Protocol (SNMP).

SNMP is a network management application layer protocol used for managing IP-based devices. The SNMP architecture requires an SNMP manager and an SNMP agent. An SNMP manager receives messages from an SNMP agent or sends messages to an agent. The messages sent between the SNMP manager and agent are termed traps. These traps contain a tree-structure that identifies occurrence and resources triggering the trap. The tree-structure is called a Management Information Base (MIB), which is an international standard for formatting SNMP messages. Today, all IP devices provide an SNMP agent.

The SNMP manager is typically a network management system such as CiscoWorks 2000, HP Openview, Sun Sunnet Manager, or the IBM Tivoli TMS10 network management system. An SNMP manager provides for

management functions through the use of traps or Inform requests. Using SNMP messages an SNMP manager can perform the following tasks:

- Set router configuration parameters
- Reload a router
- Retrieve performance data
- Receive unsolicited notification messages
- Retrieve hardware configuration information

Notification messages contain the SNMP MIB; however, a trap notification does not guarantee that the SNMP manager has received the notification message. An inform message will retry sending the message if the SNMP manager does not positively reply to the inform message with an SNMP response PDU message.

Cisco IOS supports SNMP Version 1 and SNMP Version 2 C. The differences between Cisco IOS release 12.0 and 11.3 are outlined below.

Cisco IOS Release 12.0 supports

- **SNMPv1**, the Simple Network Management Protocol version 1, as specified in IETF RFC 1157.
- **SNMPv2C**, which consists of the following:
 - **SNMPv2**: Version 2 of the Simple Network Management Protocol, an IETF draft defined in RFCs 1902 through 1907.
 - **SNMPv2C**: The Community-based Administrative Framework for SNMPv2, an experimental IP defined in RFC 1901.

Cisco IOS Release 11.3 removed support for the following version of SNMP:

- **SNMPv2Classic**: The IETF-Proposed Internet Standard of version 2 of the SNMP, defined in RFCs 1441 through 1451.

All other Cisco IOS releases support SNMPv1.

The **snmp-server community** command is used for enabling the SNMP agent function on the router. The format of the command is as follows:

```
snmp-server community string [view view-name] [ro | rw] [number]
```

The *string* variable is the access character-string used for identifying the SNMP management community that will be monitoring the router using SNMP. The optional keyword **view** and its variable *view-name* is used to restrict the SNMP information available to the SNMP managers using the community string. The **ro** keyword denotes read-only and the **rw** keyword denotes read-write capabilities for this community. The *number* optional

variable is a number relating to an IP access list that can further restrict which SNMP managers use the provided community string for obtaining information from/to the router.

The **view** keyword *view-name* variable value must match a previously defined view if it is included. The format for defining the view is

snmp-server view *view-name oid-tree (included | excluded)*

The *view-name* is a label given to the view being defined. The *oid-tree* variable is an object identifier that follows the ASN.1 subtree standard. The value can either be the subtree string of numbers, such as 1.3.6.2.4, or a word that defines a hierarchy of the ANS.1 subtree string, such as system. You can also use a wild card "*" character in the numeric string to specify a complete subtree family. The **include** **exclude** keywords are used to define the scope of the view.

For example, the command defines a view for all MIB-II system group objects with the exception of sysService (System 7):

```
snmp-server view sys system included
snmp-server view sys system.7 excluded
snmp-server view sys cisco included
```

The following commands provide information that can be obtained by an SNMP manager and are useful for troubleshooting:

snmp-server contact *text*

snmp-server location *text*

snmp-server chassis-id *number*

Using these commands, the name of the contact for supporting the router, the location of the router, and the serial number of the router can be provided and queried at any time by an SNMP manager. The *text* variable is a free-form character string that can be used for any reason to convey information to a network operator responsible for managing the router. The *number* variable of the **snmp-server chassis-id** command is typically used for the serial number of the router but can also contain the router name.

For optimal use of the SNMP agent on the router, the default traps and specific traps must be sent to a specific SNMP manager by using the **snmp-server host** command. The default traps, such as interfaces becoming active or inactive, router reloads, or configuration changes, do not need to be enabled. All other supported traps, however, must be enabled using the **snmp-server enable traps** command.

The formats of these two commands are as follows:

snmp-server host *host [version {1 | 2c}] community-string*

[udp-port port] [notification-type]

snmp-server enable traps *[notification-type] [notification-option]*

The *host* variable is the IP name of the SNMP manager or the IP address of the SNMP manager. The *community-string* variable is the community string associated with the SNMP manager. The *port* variable of the **udp-port** keyword allows you to modify the UDP port number from the default UDP port number 162. The *notification-type* variable of the **snmp-server enable traps** command indicates the specific type of trap to send to the SNMP manager. Not specifying the *notification-type* variable results in all types of traps being sent. Table 12-8 lists the possible *notification-type* variable values.

The *notification-option* is used to enable specific notification-types for the **envmon**, **isdn**, **repeater**, or **snmp** notification types that are selected. Table 12-9 lists the variable for the notification-option variable.

Table 12-8

The notification-type Values for Use with the **snmp-server Traps** and **snmp-server Informs** Commands

| <i>notification-type</i> value | Description of <i>notification-type</i> |
|--------------------------------|---|
| bgp | Sends Border Gateway Protocol (BGP) state change notifications. |
| config | Sends configuration notifications. |
| entity | Sends Entity MIB modification notifications. |
| envmon | Sends Cisco enterprise-specific environmental monitor notifications when an environmental threshold is exceeded. When the envmon keyword is used, you can specify a notification-option value. |
| frame-relay | Sends Frame Relay notifications. |
| isdn | Sends Integrated Services Digital Network (ISDN) notifications. When the isdn keyword is used on Cisco 1600 series routers, you can specify a notification-option value. |
| repeater | Sends Ethernet hub repeater notifications. When the repeater keyword is selected, you can specify a notification-option value. |
| rtr | Sends response time reporter (RTR) notifications. |
| snmp | Sends Simple Network Management Protocol (SNMP) notifications. When the snmp keyword is used, you can specify a <i>notification-option</i> value. |
| syslog | Sends error message notifications (Cisco Syslog MIB) and specifies the level of messages to be sent with the logging history level command. |

The *host* variable is the IP name of the SNMP manager or the IP address of the SNMP manager. The *community-string* variable is the community string associated with the SNMP manager. The *port* variable of the **udp-port** keyword allows you to modify the UDP port number from the default UDP port number 162. The *notification-type* variable of the **snmp-server enable traps** command indicates the specific type of trap to send to the SNMP manager. Not specifying the *notification-type* variable results in all types of traps being sent. Table 12-8 lists the possible *notification-type* variable values.

The *notification-option* is used to enable specific notification-types for the **envmon**, **isdn**, **repeater**, or **snmp** notification types that are selected. Table 12-9 lists the variable for the notification-option variable.

Table 12-8

The notification-type Values for Use with the **snmp-server Traps** and **snmp-server Informs** Commands

| <i>notification-type</i> value | Description of <i>notification-type</i> |
|--------------------------------|---|
| bgp | Sends Border Gateway Protocol (BGP) state change notifications. |
| config | Sends configuration notifications. |
| entity | Sends Entity MIB modification notifications. |
| envmon | Sends Cisco enterprise-specific environmental monitor notifications when an environmental threshold is exceeded. When the envmon keyword is used, you can specify a notification-option value. |
| frame-relay | Sends Frame Relay notifications. |
| isdn | Sends Integrated Services Digital Network (ISDN) notifications. When the isdn keyword is used on Cisco 1600 series routers, you can specify a notification-option value. |
| repeater | Sends Ethernet hub repeater notifications. When the repeater keyword is selected, you can specify a notification-option value. |
| rtr | Sends response time reporter (RTR) notifications. |
| snmp | Sends Simple Network Management Protocol (SNMP) notifications. When the snmp keyword is used, you can specify a <i>notification-option</i> value. |
| syslog | Sends error message notifications (Cisco Syslog MIB) and specifies the level of messages to be sent with the logging history level command. |

Table 12-10

The notification-type Values for Use with the **snmp-server host traps|informs** Command

| <i>notification-type values</i> | <i>Description of notification-types</i> |
|---------------------------------|--|
| bgp | Sends Border Gateway Protocol (BGP) state change notifications. |
| config | Sends configuration notifications. |
| dspu | Sends downstream physical unit (DSPU) notifications. |
| entity | Sends Entity MIB modification notifications. |
| envmon | Sends Cisco enterprise-specific environmental monitor notifications when an environmental threshold is exceeded. |
| frame-relay | Sends Frame Relay notifications. |
| isdn | Sends Integrated Services Digital Network (ISDN) notifications. |
| llc2 | Sends Logical Link Control type 2 (LLC2) notifications. |
| rptr | Sends standard repeater (hub) notifications. |
| rsrb | Sends remote source-route bridging (RSRB) notifications. |
| rtr | Sends response time reporter (RTR) notifications. |
| sdlc tions. | Sends Synchronous Data Link Control (SDLC) notifications. |
| sdllc | Sends SDLLC notifications. |
| snmp | Sends Simple Network Management Protocol (SNMP) notifications defined in RFC 1157. |
| stun | Sends serial tunnel (STUN) notifications. |
| syslog | Sends error message notifications (Cisco Syslog MIB). The level of messages to be sent can be specified with the logging history level command. |
| tty | Sends Cisco enterprise-specific notifications when a Transmission Control Protocol (TCP) connection closes. |
| x25 | Sends X.25 event notifications. |

Because inform requests provide a reliable delivery of the notification message, the resending parameters can be defined to meet the networks expectation. The format of specifying the resend operational parameters for

Table 12-10

The notification-type Values for Use with the **snmp-server host traps|informs** Command

| <i>notification-type values</i> | <i>Description of notification-types</i> |
|---------------------------------|--|
| bgp | Sends Border Gateway Protocol (BGP) state change notifications. |
| config | Sends configuration notifications. |
| dspu | Sends downstream physical unit (DSPU) notifications. |
| entity | Sends Entity MIB modification notifications. |
| envmon | Sends Cisco enterprise-specific environmental monitor notifications when an environmental threshold is exceeded. |
| frame-relay | Sends Frame Relay notifications. |
| isdn | Sends Integrated Services Digital Network (ISDN) notifications. |
| llc2 | Sends Logical Link Control type 2 (LLC2) notifications. |
| rptr | Sends standard repeater (hub) notifications. |
| rsrb | Sends remote source-route bridging (RSRB) notifications. |
| rtr | Sends response time reporter (RTR) notifications. |
| sdlc tions. | Sends Synchronous Data Link Control (SDLC) notifications. |
| sdllc | Sends SDLLC notifications. |
| snmp | Sends Simple Network Management Protocol (SNMP) notifications defined in RFC 1157. |
| stun | Sends serial tunnel (STUN) notifications. |
| syslog | Sends error message notifications (Cisco Syslog MIB). The level of messages to be sent can be specified with the logging history level command. |
| tty | Sends Cisco enterprise-specific notifications when a Transmission Control Protocol (TCP) connection closes. |
| x25 | Sends X.25 event notifications. |

Because inform requests provide a reliable delivery of the notification message, the resending parameters can be defined to meet the networks expectation. The format of specifying the resend operational parameters for

Table 12-9

The notification-type Values for Use with the notification-option Variable for the **snmp-server traps** and **snmp-server informs** Commands

| <i>notification-type</i> Values | Possible <i>notification-option</i> |
|---------------------------------|--|
| envmon | voltage shutdown supply fan temperature |
| isdn | call-information —enables an SNMP ISDN call-information notification for the ISDN MIB subsystem. isdnu-interface —enables an SNMP ISDN U interface notification for the ISDN U interface MIB subsystem. |
| repeater | health —enables IETF Repeater Hub MIB (RFC 1516) health notification. reset —enables IETF Repeater Hub MIB (RFC 1516) reset notification. |
| snmp | authentication —enables SNMP Authentication Failure notifications. |

>>>If no notification-option is coded, then all traps/informs are sent for the notification-type.<<<

You can stop traps from being sent on interfaces that are expected to become active or inactive by using the **no snmp trap link-status** interface command. For instance, in order to stop traps/informs on ISDN dial-up connections, the following can be entered on the ISDN interface definitions:

```
interface BRI0
  no snmp-trap link-status
```

Using this configuration on the ISDN interface definition, the SNMP agent does not send traps/informs to the SNMP manager when the ISDN line becomes active and goes inactive.

For a more reliable method of sending notification messages to an SNMP message, the **snmp-server host** command is used as

```
snmp-server host host informs [version {1 | 2c}] community-string
[udp-port port] [notification-type]
```

The parameters are used in the same fashion as those specified for the trap version of the command. The **snmp-server enable traps** command with any specified *notification-type* and *notification-option* specified is also required to denote the types of inform requests sent to the SNMP manager.

The *notification-type* on the **snmp-server host traps|informs** command are listed in Table 12-10.

using inform request messages is

```
snmp-server informs [retries retries] [timeout seconds]
[pending pending]
```

The *retries* variable is the number of times the inform request process used by this SNMP agent will attempt to send the notification message to the SNMP manager before aborting the send process. The *seconds* variable is the amount of time the SNMP agent will wait between retries. The *pending* variable is the total number of outstanding acknowledgments of inform requests sent to the SNMP manager. Once the *pending* value is reached, the oldest outstanding unacknowledged inform requests are discarded.

A final SNMP command useful on the router is the **snmp-server trap-source** command. The format of this command is **snmp-server trap-source interface**

The *interface* variable is the name of any active interface on the router. Using a defined interface, loopback 0 is an optimal choice to be the source address for the SNMP agent since it is always considered to be active.



NOTE: *SNMP managers use a basic IP PING command to determine if a router is reachable on the network. Using the IP address of the interface loopback 0 virtual interface enables the SNMP manager PING to reach the router on any available active network interface on the router.*

An example of enabling an SNMP agent on a router is as follows:

```
snmp-server community allmgrs
snmp-server trap-source loopback0
snmp-server chassis-id 7000-234-5654
snmp-server contact Network Operations
snmp-server location 1770 Orchard Lane, Atlanta, GA
snmp-server enable traps
snmp-server host Prismsnp netmgrs snmp
snmp-server host Secsnmp netmgrs snmp
```

In this example, the SNMP community string used by all SNMP managers for reading or writing to this router must use the allmgrs string. All traps and inform requests sent from this router will have the IP address of interface loopback0. Two SNMP managers, Prismsnp and Secsnmp, will be receiving all SNMP traps and inform requests.

Using the Banner Command

A final command that can be useful for displaying information to the router administrators is the banner command. It enables the display of information at various entrance points to the establishment of a successful connection:

```

banner exec ^
*****
*
*   ROUTER NAME: c1r2LNDN SN 77778888           *
*   LOCATION:   10 Downing St.                 *
*              Newark, NJ                     *
*   CONTACT:   Sam Samson                       *
*              201-222-2222                     *
*   CONTACT PAGE: 800-888-8888                 *
*   Cisco TAC:  800-553-2447                   *
*   Cisco ACCT: 1111111                       *
*****
banner motd ^
WARNING: The unauthorized use of devices or the tampering with
access
           to this network and its resources are CRIMINAL offenses.
           *** ALL VIOLATORS ARE SUBJECT TO PROSECUTION *** ^C
    
```

In this example, the **banner** command is used for displaying a warning message for unauthorized use of the router. This is done using the **motd** (message of the day) keyword of the **banner** command. The ^ character marks the beginning and end of the text to be displayed. The **motd** banner is displayed before any operational router access is attempted, including anything prior to the login password prompt.

The **banner exec** command displays the banner once the router administrator has gained access to the EXEC mode on the router. We use this banner for our example to denote the location of the router, contact phone numbers, and the Cisco Systems maintenance contract code along the router serial number.

The other points of banner displays are listed in Table 12-11.


Table 12-11
The Banner Command Display Points Parameters

| Banner Display Point | Description of Point |
|----------------------|------------------------------------|
| exec | Sets EXEC process creation banner |
| incoming | Sets incoming terminal line banner |
| login | Sets login banner |
| motd | Sets message of the day banner |

CHAPTER

13

IP Configuration



The basis for building routed networks lies in defining IP subnets to the interfaces on a router that connect LANs and WANs. This chapter focuses on the definition of IP protocol and service provided by the implementation of the IP protocol on Cisco routers.

Defining Subnets on the Router Interfaces

The definition of IP subnets on a router interface enables that interface to transport IP-based application traffic. Defining the IP addresses on an interface requires a planned implementation of the IP subnet space. The space refers to the projected number of networks and hosts needed to be supported on the corporate-wide network. For each physical interface on a router, a new IP network or IP subnet is required, which is a basic IP network-addressing requirement. If a network has 100 individual LANs, for instance, then it is possible that there is a need for a minimum of 100 networks or subnets. Since it is not possible for any one corporation to attain 100 IP networks from the Internet Assigned Numbers Authority (IANA), a single IP network is used and this network is subnetted.

Within each subnet, the number of supported hosts must be determined. A router interface is represented as a host on the IP subnet being defined. Table 13-1 reviews the list of IP address class and the available range of IP networks within each class.

In addition to the IP networks identified in Table 13-1, the IETF has reserved address space for any network to use. The requirement is that the IP addresses are not advertised on the Internet. This is because the following addresses are not registered with the IANA and therefore are not unique, which creates the potential of routing problems. These “open” IP network addresses are defined in the Internet Engineering Task Force’s (IETF) RFC 1918, which describes their general use for any IP network.

In planning your IP address schema, it is prudent to define a network addressing standard. For instance, the host address of the IP address assigned to an interface is always 1, IP addresses assigned to Windows NT servers fall within a range of consecutive host numbers, and send user workstations fall within a range of consecutive host numbers. Diligence in planning the IP address allocation facilitates the ease of assignment and simplifies initial troubleshooting analysis.

Table 13-1
IP Address Class
Address Ranges
and Their Associ-
ated Statuses

| IP Address Class | IP Address or Range | Status |
|------------------|------------------------------|---------------------------|
| A | 0.0.0.0 | Reserved |
| | 1.0.0.0 to 126.0.0.0 | Available |
| | 127.0.0.0 | Reserved |
| B | 128.0.0.0 to 191.254.0.0 | Available |
| | 191.255.0.0 | Reserved |
| C | 192.0.0.0 | Reserved |
| | 192.0.1.0 to 223.255.254 | Available |
| | 223.255.255.0 | Reserved |
| D | 224.0.0.0 to 239.255.255.255 | Multicast group addresses |
| E | 240.0.0.0 to 255.255.255.254 | Reserved |
| | 255.255.255.255 | Broadcast |

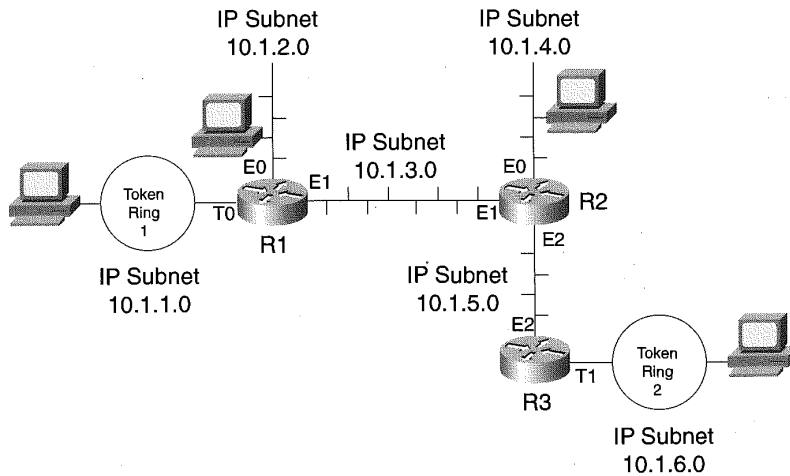
Table 13-2
RFC 1918 Non-reg-
istered IP Network
Address Ranges for
Any IP Network

| RFC 1918 IP Address Class | RFC 1918 IP Address or Range | Status |
|------------------------------|---------------------------------|---|
| A | 10.0.0.0 | The entire address space available for use. |
| B | 172.16.0.0–172.31.0.0 | The entire address space available for use. |
| C | 192.168.0.0 | The entire address space available for use. |

Assigning an IP address to a router interface is performed while under interface configuration mode. Recall that to enter configuration mode the router administrator must first access the privileged command line interface (CLI) function of the router. Once in configuration mode, the router administrator enters the name of the interface type to which the IP address is being applied. Using Figure 13-1 as an example, we can assign IP addresses to the routers and interfaces shown.

Figure 13-1

Sample network configuration for assigning IP addresses to router LAN interfaces.



Assume that, outside of the configuration shown, each router is reachable over the corporate network. The router administrator Telnets to the already active IP interface on each router. Once in each router, the router administrator access configuration mode. The procedure for this on each router is

```

routername> enable
password: xxxxxxxx
routername#> configuration terminal
routername(config)#>interface ethernet 0

```

At this point, the router administrator enters the interface type (interface Ethernet 0) being given an IP address. Assigning the IP address is performed by issuing the following command:

ip address *ip-address mask*


The *ip-address* positional variable value is the complete IP address being assigned to the interface. The positional *mask* variable value defines the bits used against the address to determine the IP subnet and host value used to access the interface with IP. Cisco IOS only supports masks using contiguous bits flushed left in the addressing. For example, a mask of 255.255.255.0 against a Class B address is represented in bit notation as

| Class B Network number | subnetwork | host number |
|-------------------------|-------------|-------------|
| 1111 1111 . 1111 1111 . | 1111 1111 . | 0000 0000 |

If this example is based on a Class B address, then the third octet is masking the value as a subnet of the network. For instance, 172.16.8.0 is a different subnet than 172.16.9.0 using the above mask. If we applied a network mask using the following decimal representation of 255.255.252.0, however, the mask is as follows:

| Class B Network | subnet | host number |
|-------------------------|---------|----------------|
| 1111 1111 . 1111 1111 . | 1111 11 | 00 . 0000 0000 |

If this is the applied mask to the sample 172.16.8.0 subnet, then the 172.16.9.0 addressing is also found in the same subnet. The use of the consecutive bits to the left in the mask has increased the number of host addresses and decreased the number of possible networks. In this example, the actual complete subnet range is 172.16.8.0–172.16.11.255. This subnetting provides 1,022 host addresses and 62 subnet addresses. This subnetting range follows IETF RFC 950. Cisco IOS, however, enables the full range of an IP network to be used for addressing, as we will see further on, by allowing the use of subnet 0 within an address space.



NOTE: An IP subnet can be defined to only one active router interface at any given time. Router interfaces in shutdown (inactive) mode can have the same subnet or, for that matter, the same IP address assigned. It is strongly suggested to not do this, however, as it can lead to confusion during troubleshooting.

Using the following configuration commands, the routers shown in Figure 13-1 are assigned IP addresses. The Ethernet interfaces are represented in the figure using an E and the Token Ring interface are represented in the figure using a T. The number following each interface is the port used on the interface processor for connecting to the network.

Router R1 Configuration:

```
interface Ethernet 0
ip address 10.1.2.1 255.255.255.0
interface Ethernet 1
ip address 10.1.3.1 255.255.255.0
interface Tokenring 0
ip address 10.1.1.1 255.255.255.0
end
```

Router R2 Configuration:

```
interface Ethernet 0
ip address 10.1.4.1 255.255.255.0
interface Ethernet 1
ip address 10.1.3.2 255.255.255.0
```

```
interface Ethernet 2
ip address 10.1.5.1 255.255.255.0
end
```

Router 3 Configuration:

```
interface Ethernet 2
ip address 10.1.5.2 255.255.255.0
interface Tokenring 1
ip address 10.1.6.1 255.255.255.0
end
```

The IP addresses assigned in the sample configuration are referred to as the primary IP address for each interface. In the sample configurations, we see that the mask 255.255.255.0 is applied to a Class A IP address of 10.0.0.0. The subnet mask is actually 0.255.255.0 since the IP address is a Class A network address. On the interface configuration command **ip address**, however, the *mask* variable value is expressed as all the bits used for denoting the network and the subnet of the IP address assigned. A mask of this nature is also referred to as a 24-bit mask.

In the sample configuration, the Ethernet LAN attaching router R1 to R2 requires that the interfaces on each of these routers use a unique host number in the IP address. The same requirement goes for the definitions of the Ethernet interfaces connecting R2 with R3.

When to Assign Multiple IP Addresses to an Interface

Cisco IOS allows for multiple IP addresses on a single router interface. The IP addresses assigned to an interface that are not the primary router interface are termed secondary IP addresses. There is no restriction on the number of secondary IP addresses that can be defined to a router interface. Typically, multiple IP addresses are assigned to a router interface in one of the following situations:

- Host address allocation has become insufficient to support the number of IP addressed devices on a single physical network.
- For supporting the migration from a bridged IP network to a routed IP network.
- For connecting two subnets of a single network that are currently connected through another network.

Suppose the network has the topology shown in Figure 13-2. Here we see that on router R1 Ethernet 0 all 253 host addresses have been allocated. Seven new workstations are being added to the Ethernet segment to support the expansion of the department. The router administrator can accommodate the expansion by adding a secondary IP address that specifies another subnet to be associated with Ethernet 0. The router administrator performs the following commands in configuration mode on router R1:

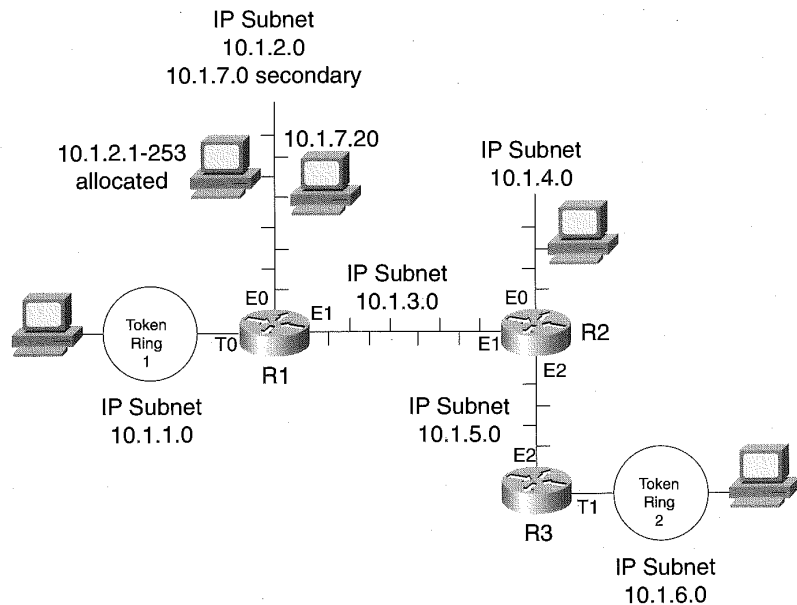
Router R1 Configuration:

```
interface Ethernet 0
ip address 10.1.7.1 255.255.255.0 secondary
end
```

The router administrator has added a secondary subnet to Ethernet 0 on router R1 that provides another 253 possible host addresses for connecting IP hosts to the network. It is the addition of the **secondary** keyword to the **ip address** command that denotes as additional IP address assignment on the interface being configured. The resulting configuration for Ethernet 0 on router R1 looks like the following:


Router R1 Configuration:

Figure 13-2
Using multiple IP addresses on router interfaces.



```
interface Ethernet 0
ip address 10.1.2.1 255.255.255.0
ip address 10.1.7.1 255.255.255.0 secondary
interface Ethernet 1
ip address 10.1.3.1 255.255.255.0
interface Tokenring 0
ip address 10.1.1.1 255.255.255.0
```

From the configuration sample for router R1, we can see that the configuration for Ethernet 0 uses 10.1.2.1 as the primary IP address and 10.1.7.1 becomes the secondary IP address, thus providing a second subnet to the Ethernet LAN. Now the LAN, assuming traffic load does not cause performance problems, can support up to 506 connected IP devices through the Ethernet 0 interface of router R1.



NOTE: All routers connecting to the same LAN must define the same IP subnet address in order to participate in delivering IP traffic for the LAN. Each router on the same LAN, however, must use a different host number within the same IP subnet. Try using Hot Standby Routing Protocol (HSRP) to allow multiple routers the capability to respond to the same IP address.

Maximizing Network Address Space Using Subnet Zero

The total number of networks available in any given subnet is the decimal value of the bits used in defining the mask for the IP address range minus one. For example, a Class B address (172.16.0.0) with a mask of 255.255.255.0 (0.0.255.0) has the potential for 255 subnets. However, the all-zeroes subnet, 172.16.0.0, is reserved in the RFC 791 specifications to denote the network address. The all-ones subnet, 172.16.255.0, though it can be explicitly defined, is discouraged from use. The format for specifying the use of subnet zero as a viable addressable subnet is as follows:

ip subnet-zero

This command is a global command and enables the use of subnet zero on all the routers' interfaces. The feature is disabled if not specified. Be aware that if employing the subnet zero feature, the 172.16.0.0 network address in our example, subnetted using 255.255.255.0, is written exactly the same and hence is distributed by the routers as both a network and a subnet.

Routing Same Network Packets Without a Routing Table Entry

At times a router may encounter a packet that is to be sent to a subnet within its own network, but the routing table within the router does not have a route to the destination subnet. In such a case, the router throws the packet of the unknown destination subnet away. Since the destination subnet is within the major network addressing of the overall IP network, however, a route should be available to the destination. Cisco IOS provides such a method using IP classless routing.

For example, in Figure 13-3, a workstation sends a packet with 10.8.1.0 as the destination. The local router receives the packet and inspects its routing table. In doing so, the router determines that a specific route to the destination does not exist. Using the IP classless routing feature, however, the router sends the packet out an interface that meets the network addressing portion of the destination. Specifying the global configuration command **ip classless** under configuration mode enables this feature. Once this command is specified, it applies to all routing protocols used by the router.

The following lists some instances when the classless feature is used:

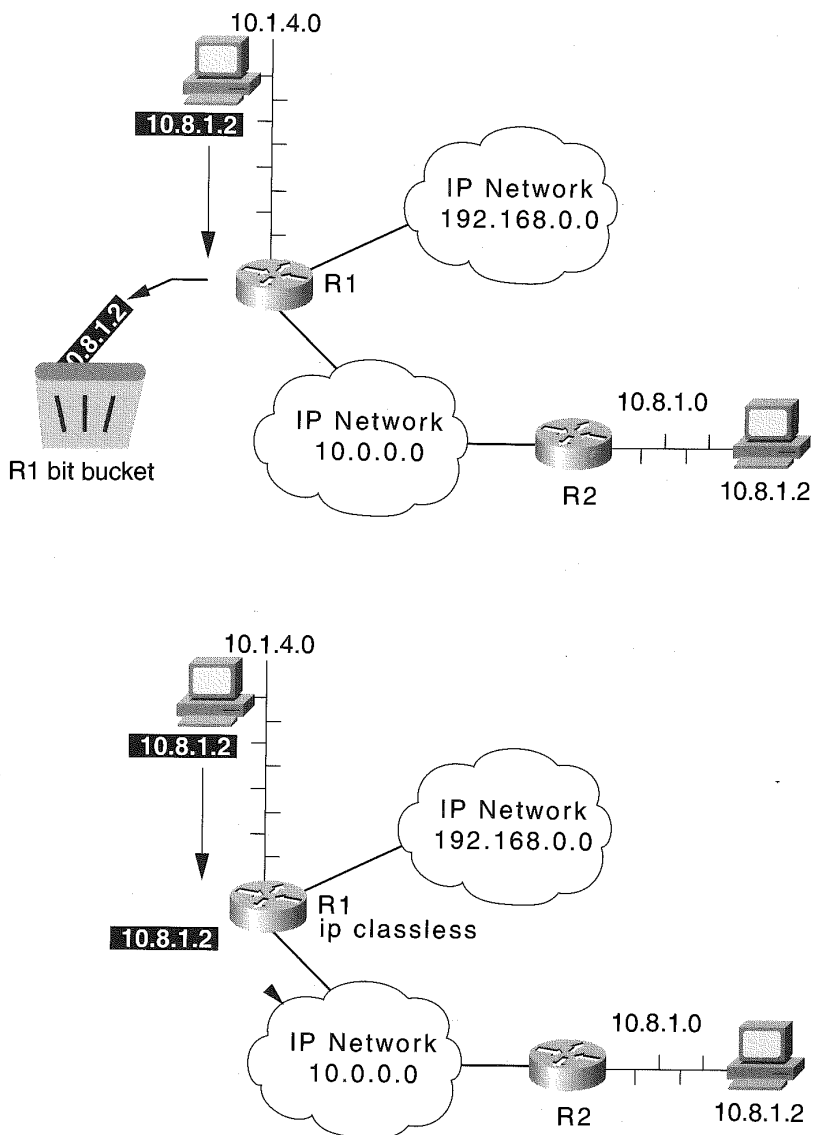
- For multiple routing protocols within a network.
- When using the passive keyword on a router's only interface to the WAN.
- When static addressing is used on a router, instead of a routing protocol.

Enable IP on a Serial Interface Without Specifying an IP Address

Many designers choose to save the IP address space by not assigning an IP address to the WAN serial interfaces. The reasoning behind this is that many subnets are utilized for two host addresses: one serial interface on one router and the serial interface on another router. This means that a subnet supporting 254 host addresses wastes 252 addresses needlessly. To avoid the waste of addresses but maintain the capability of the serial interface to carry IP traffic, the Cisco IOS employs a feature called IP unnumbered.

Figure 13-3

Use of the IP classless feature to route unknown subnets to the best supernet interface.



The command for specifying IP unnumbered has the following format:
ip unnumbered *type number*

The *type* positional variable value is an active interface on the router with a valid IP address. The *number* variable denotes the specific port on the interface type being used as the source IP address of the packets, leaving the router over the serial interface.

The command is specified under the interface configuration mode for the serial interface being assigned the unnumbered feature. The following is an example of such a configuration:

```
interface loopback 0
ip address 10.10.10.1 255.255.255.0
interface serial 4/1
ip unnumbered loopback 0
```

Using this configuration, the router sends all packets from interface serial 4/1 with the IP address 10.10.10.1 as the source of the packet. This technique can save IP address space. The command can be applied to HDLC, PPP, LAPB, Frame Relay, ATM, SLIP, and tunnel interfaces. It cannot be applied to SMDS or X.25 interfaces.

Because the IP address is “borrowed” from another active interface, an SNMP manager cannot determine if the interface is up or down by pinging the borrowed IP address. The SNMP manager, however, can receive traps or informs from the router on the status of the interface. Likewise, you cannot ping the actual interface, perform a netboot over the interface, or support IP security options on an unnumbered serial interface.

In the configuration example, the loopback interface was chosen since the router always interprets this interface as being active and connected to the network.

IP Address Mapping

Part of the role of the router is to resolve IP addresses. The router resolves IP addresses through a mapping process using various protocols. These are:

Address Resolution Protocol (ARP): Known IP-address mapping to a 48-bit MAC address


Proxy ARP: A process in which the router places its own MAC address in the outbound packet

Reverse ARP: Known 48-bit MAC-address mapping to an IP address

ARP is a request sent by an IP host to determine the Layer 2 MAC address of the destination IP host device. Seen in Figure 13-4, the requesting ARP is sent into the network using the MAC broadcast address

Figure 13-4

ARP process for determining the MAC address of the destination IP host.



0xFFFFFFFFFFFF'. Each IP device receives the ARP packet and inspects the packet as to whether the target IP address specified is indeed its own IP address. If the IP address does not match, the IP host discards the packet. If the IP address does match, the matching IP host places its own MAC address in the target MAC field of the ARP reply packet and sends the ARP reply back to the original requesting IP host.

Routers assist an IP host in mapping the target IP address to a MAC address by using proxy ARP. Shown in Figure 13-5, the proxy ARP function of the router replies back to the requesting ARP workstation on behalf of the actual destination IP host. This is possible because of the router building an ARP table based on all the packets that traverse the interfaces of the router and the use of the routing table.

Upon receiving the ARP request, the router determines if it has IP address to MAC address mapping already completed for the target IP address in the ARP. This is done by searching the ARP table (ARP cache) built by the router. If it finds a match of the IP address and has a valid route to the subnet of the matching IP address, the router performs the reply but uses its own MAC address, instead of the MAC address of the true target device.

The proxy ARP function is enabled by default on a Cisco router. It is on by default because any time an IP host tries to communicate with another IP host that is not in its local ARP cache, the origin IP host must send an ARP. Having the Cisco router act as an ARP server using proxy ARP functionality eliminates unproductive packet delivery throughout the network.

Reverse ARP, as illustrated in Figure 13-6, works in the same manner as ARP. The difference is the discovery of the IP address of the requesting host, instead of the MAC address of a target IP host. A RARP server is required on the network to perform this function. Since the purpose of RARP is to provide an IP address for the requesting IP host, it is often associated with diskless workstations that normally do not store their network configuration between power-on resets. Cisco IOS can be configured on a router to

Figure 13-5

The proxy ARP process for reducing unproductive packets to resolve MAC addresses of destination IP hosts.

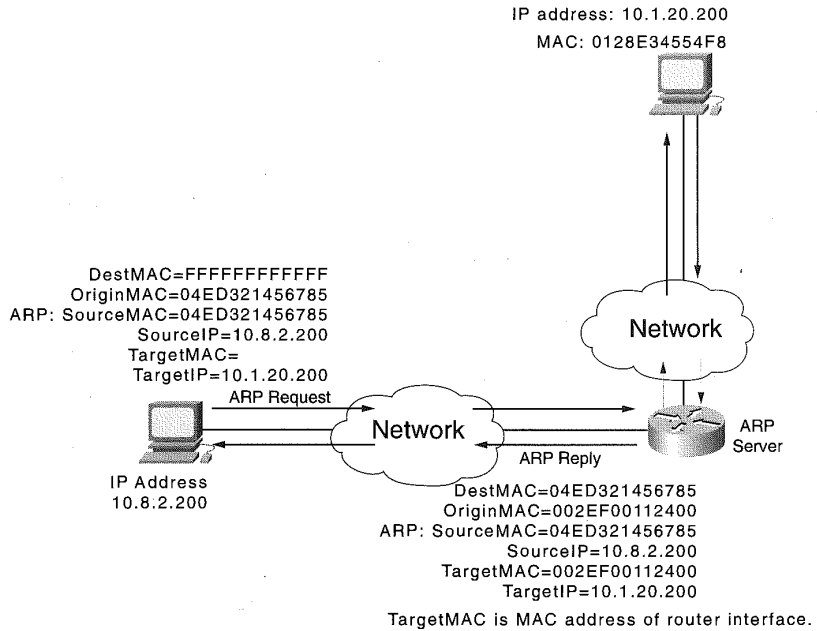
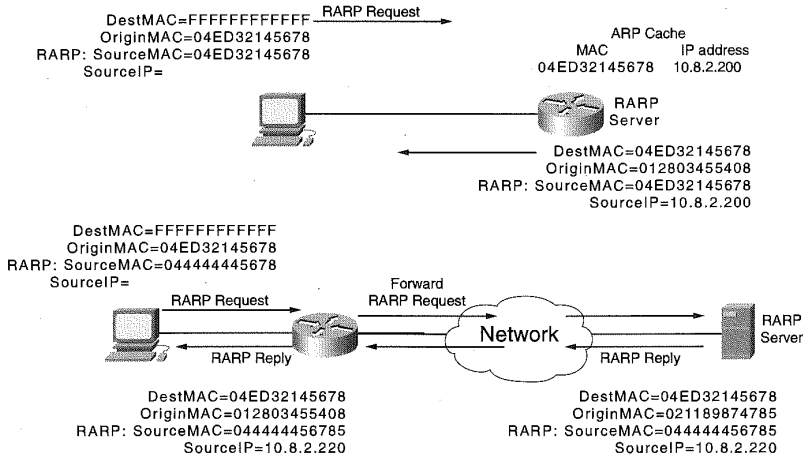


Figure 13-6

RARP process for determining IP address of a host.



support a RARP server function for RARP requests when a static IP to MAC address mapping has been defined for the requesting device. The router can also forward the RARP requests directly to a known RARP server.

The dynamic ARP cache entries built by the router have a default timeout of four hours. The timeout ARP cache entries can be changed by using the global or interface configuration command

arp timeout *seconds*

where *seconds* can be from 0 to 14400. A value of 0 indicates that the entries never have a timeout.

Define a Static ARP Cache

Routers perform ARP caching by default, but in some cases a static ARP entry may be needed, such as in supplying an IP address for a RARP request from a diskless workstation. In this case, the following global configuration command must be entered to define static ARP cache entries:

arp *ip-address hardware-address type[alias]*

The **arp** global command allows the router administrator to map an IP address (*ip-address*) of a host that is found through any directly attached interface defined on the router with a MAC address (*hardware-address*). The *type* variable identifies the encapsulation type being used for the mapping. The type can be **arpa**, **snap**, or **probe**. The **arpa** value is the default and supports all Ethernet resources, **snap** is used for FDDI and Token Ring resources, and **probe** is a Hewlett-Packard-specific protocol (HP-Probe) that can be used on any interface type for resources supporting the HP-Probe protocol. The HP-Probe is not greatly used. The *alias* keyword directs the Cisco IOS to respond to the ARP request with its own hardware address as though the router were the owner of the IP address.

For example, to define an ARP cache entry for the diskless workstation in Figure 13-6, the following definition was provided on the router:

```
arp 10.8.2.200 04ED.3214.5678
interface ethernet 0
ip address 10.8.2.1 255.255.255.0
ip rarp-server 10.8.2.1
```

Using this definition, the *type* variable of the **arp** command is defaulting to **arpa**. This indicates to us that the workstation using the 04ED.3214.5678 MAC address resides on an Ethernet LAN segment. Using this entry, the router replies back with the 10.8.2.200 IP address to an RARP request from the 04ED.3214.5678 MAC address. The **ip rarp-server** interface configuration command enables the router to act as a RARP server. The **ip rarp-server** command identifies the 10.8.2.1 IP address as the source address in the RARP reply.

Supporting Multiple ARP Encapsulation Types on an Interface

The ARP encapsulation type supported on any interface can be modified to support more than just the default arapa encapsulation. When working with the **arp** interface configuration command, one or more encapsulation techniques can be used. The following describes the three different types supported and how they are specified for the interface:

arp arpa: IEEE 802.3 Ethernet, as described with RFC 826 (the default value)

arp probe: HP-Probe protocol for IEEE 802.3 networks

arp snap: ARP packets in support of RFC 1402 for FDDI and Token Ring networks

Disable Proxy ARP

Cisco IOS has a proxy ARP function enabled by default. Disabling it forces the router to deliver the Layer 2 broadcast packet along the supported route for the IP subnet specified in the ARP request. Entering the following global configuration command disables the proxy ARP function:

```
no ip proxy-arp
```

To reenble the proxy ARP function after it has been disabled, enter the global configuration command

```
ip proxy-arp
```



NOTE: All Cisco IOS commands have their definitions or functions negated by entering a “no” in front of the enabling command under configuration mode.

Enabling the Use of IP Host Names in Cisco IOS Commands

Cisco IOS has several commands that support the use of a host name within the command, instead of an IP address. Examples of the commands are connect, telnet, ping, write network, and configuration network. The IP

address to IP host name mapping within the Cisco router is enabled by defining specific IP host names and their IP addresses and/or by implementing Domain Naming System (DNS) services.

Static Definition of IP Host Names and IP Addresses

Manual definition of an IP address to a IP host name enters a static mapping in the host name table of the router. The format of the global configuration command for this static mapping is

```
ip host name [tcp-port-number] address1 [address2...address8]
```

The *name* positional variable is the name assigned by this router for the IP address (*address1*) defined on the command. The optional *tcp-port-number* positional variable is the TCP port used for connecting to the IP host being defined. The optional variables *address2...address8* provide up to a total of eight IP addresses that are bound to the host name defined with the command.



NOTE: The DNS standards allow multiple host names to an IP address but not multiple IP addresses to a single host name.

Using this command, suppose a server has four Ethernet attachments on four different subnets. Using the `ip host` command, we can define the router to four IP addresses that attempt to connect to the desired host. For example, in the following code

```
ip host backbone-router 10.1.1.1 10.2.2.2 10.3.3.3 10.4.4.4
```

the router attempts to connect to the IP host names `backbone-router` by trying each address in sequence. Using static IP host to address mapping provides for a fast IP address resolution.

Enable the Use of DNS Services

Cisco IOS can try to resolve IP host names by using DNS services. Entering an IP host name for a remote host can be resolved if the following global configuration command has been entered on the router:

ip domain-lookup

Having domain lookup enabled allows the Cisco IOS to query the network DNS servers for an IP address. The router sends the fully qualified name and the Cisco IOS performs the name qualification based on the use of the following global configuration command:

ip domain-name *name*

The *name* variable is the domain name used by this router to resolve the host-to-address mapping. For example, entering


```
ip domain-name networkcorp.com
```

qualifies the host name of corpmail in a ping command to corpmail.networkcorp.com as the qualified name presented to the DNS server for IP address resolution. Multiple domain names can be entered by using multiple **ip domain-list** commands to assist in resolving host names with fully qualified alias domain names. The **ip domain-list** command has the same command format as the **ip domain-name** command. The domain name(s) specified with the **ip domain-list** take precedence, in sequential order, over the default domain name defined by the use of the **ip domain-name** command.

The DNS servers used to resolve the queries are defined using the **ip name-server** global configuration command. For example, entering

```
ip name-server 10.10.10.200 10.10.20.200 10.10.30.200
```

causes the router to send DNS queries to the three IP addresses specified on the **ip name-server** command. The first response received is the IP address used for the router command that caused the query. The **ip name-server** command supports up to six DNS server addresses.



NOTE: Although the DNS server's IP addresses can be specified on a single **ip name-server** command, they are written in the configuration file as separate **ip name-server** commands. For instance, the example shown is written as

```
ip name-server 10.10.10.200
ip name-server 10.10.20.200
ip name-server 10.10.30.200
```

To complete the DNS services example, the configuration file is written as

```
ip domain-lookup
ip domain-name networkcorp.com
ip name-server 10.10.10.200
ip name-server 10.10.20.200
ip name-server 10.10.30.200
```

Disable IP Routing


Cisco IOS software enables IP routing once an IP address is specified on a router interface and an IP routing protocol is defined for use by the router. If bridging IP is a requirement for connecting the router to a network, IP routing must be disabled. This affects the entire operation of the router. The **no ip routing** global command disables the routing function on the entire router. Disabling IP routing causes the router to act as an IP end host, and it can be reenabled by specifying the **ip routing** global configuration command.

Routing Assistance When IP Routing Is Disabled

With IP routing is disabled, Cisco IOS can still learn of IP routes to other networks acting as IP host computers. The three methods used by a Cisco router to learn IP routes when IP routing is disabled are

- Proxy ARP
- Defining a default gateway
- ICMP Router Discovery Protocol (IRDP)

As described earlier, the Cisco IOS software enables the proxy ARP function by default. Using just proxy ARP, a router can provide a route to an IP host based on the ARP cache table. As a request enters the router, the router searches the ARP cache for a matching IP address and, when found, it forwards the packet on the interface.



NOTE: *If IP routing is disabled and the router is now relying on the proxy ARP function for routing packets, it is beneficial to specify a shorter ARP cache entry timeout on the router interface definitions to ensure the validity of the route.*

The second method of defining a default-gateway (router) IP address specifies to the Cisco IOS software that all packets with IP networks that are not defined on any interface specific to the router will forward the non-local packet to a specified IP address. Using the following global configuration command, a default gateway address is defined for a Cisco router:

ip default-gateway ip-address

The *ip-address* variable of the **ip default-gateway** global configuration command must be the IP address of a router that has a routing table and is attached to a network known by the router being defined. Cisco IOS supports only one ip default-gateway definition. For example, by specifying


```
no ip routing
ip default-gateway 10.1.1.2
interface Ethernet 0
ip address 10.1.1.1 255.255.255.0
```

the default gateway router is reachable via the Ethernet 0 interface of the router. Any packets received by the router from attached networks with destination networks not directly connected to the router are forwarded on the Ethernet 0 interface to the IP address 10.1.1.2 for routing of the packet. The router specified as the default gateway can either route the packet based on its routing table or reply back to the original router with an IP Control Message Protocol (ICMP) redirect message that indicates another attached router can provide a better path to the destination.

The router with IP routing disabled then caches the redirected route and uses this route for future requests to the cached IP address. This process is depicted in Figure 13-7. Using the default-gateway works is not always reliable, however, since the status of the default-gateway router is not available.

The final method employs RFC 1256 ICMP Router Discovery Protocol (IRDP). The Cisco IOS software dynamically learns of available potential default gateways by employing IRDP. It is the preferred method of the three

Figure 13-7
Using default-gateway and ICMP redirect messages for routing packets on a router with IP routing disabled.




methods discussed because IRDP can determine the availability of a default gateway and enable multiple default gateways. Cisco's implementation of IRDP is enhanced to "wire tap" an interface for RIP and IGRP routing updates. The Cisco IOS software adds the sending IP address of a RIP- or IGRP-routing update packet to the IRDP eligible default-gateway cache. Using this in addition to the sending and receiving of IRDP packets, a Cisco router can route to non-local networks with great success. Enabling IRDP is made possible by including the following interface configuration command:

```
ip irdp [multicast | holdtime seconds | maxadvertinterval seconds |
minadvertinterval seconds | preference number |
address address [number]]
```

Specifying **ip irdp** for an interface causes the Cisco IOS to broadcast IRDP advertisements with a maximum interval between advertisements of 600 seconds and a minimum advertisement interval of (600 x .075) seconds. Specifying **multicast** enables IRDP to be used with Sun Solaris systems, which respond to the multicast IP address of 224.0.0.1, instead of the all ones, or 255.255.255.255, IP broadcast address.

Using **holdtime** on the **ip irdp** command defines the length of time an IRDP advertisement is considered valid. After the **holdtime** expires, the advertisement is flushed from the IRDP table. The **holdtime** default is three times the **maxadvertinterval** value. If specified, **holdtime** must be greater than **maxadvertinterval** up to 9,000 seconds. The **maxadvertinterval** is the maximum number of seconds between sending IRDP advertisements; its default is 600 seconds. The **minadvertinterval** default is 0.75 times the **maxadvertinterval** value.



NOTE: The **holdtime** and **minadvertinterval** values will default based on the value of **maxadvertinterval**. Therefore, it is recommended that the **maxadvertinterval** be changed prior to modifying the **holdtime** or **minadvertinterval** values.

The **preference** keyword variable *number* value defines the use of this router's interface as a preferred default gateway for any other IRDP router receiving the IRDP advertisement from this router. The default value is 0 and can range anywhere from -2^{31} to 2^{31} . The higher the value, the greater the preference.

The **address** keyword allows the router administrator to specify the IP address of a known IRDP router and a *preference* value for using that router within its IRDP advertisement. In this manner, a preferred IRDP-determined default gateway can be used for all IRDP routers.

In the following example, interface Ethernet 0 enables the use of IRDP, sets its own preference to a value of 50, and advertises a second IRDP router address of 10.1.1.95 with a higher preference value of 100 as a preferred default gateway.

```
no ip routing
interface ethernet 0
  ip address 10.1.1.90 255.255.255.0
  ip irdp
  ip irdp preference 50
  ip irdp address 10.1.1.95 100
```

Using any of the above methods, or a combination of them, enables the routing of IP packets when IP routing has been disabled.

Bridging IP, Instead of Routing IP

IP is based on using a routing protocol; however, it can be bridged as well as routed at the same time. To transparently bridge IP and disable the routing of IP, the following commands must be entered into the router configuration:

```
no ip routing
interface type number
  bridge-group group
  bridge number protocol [ieee | dec]
```

Specifying the global configuration command, **no ip routing**, disables the routing of IP traffic on the router. The interface configuration command directs the configuration mode to the specified interface *type* and port *number* or slot/port *number* of the interface being configured.

The **bridge-group** *group* interface command defines the common bridge-group number applied to the interface for building a transparent bridge network. The global configuration command **bridge protocol** is required to identify the type of spanning tree protocol in use for a specific bridge-group. The number variable of the bridge protocol command ranges from 1 to 63 and identifies the bridge-group number that will use the specified spanning tree protocol. The spanning tree protocol available is either **ieee** or **dec**. For example,

```
no ip routing
!
interface Ethernet 0
  ip address 10.1.1.1 255.255.255.0
  bridge-group 1
!
```

```
interface Ethernet 1
ip address 10.1.1.2 255.255.255.0
bridge-group 1
!
bridge 1 protocol ieee
```

The value of the bridge-group command can be anywhere from 1 to 63. In the example, IP and all other protocol traffic between Ethernet 0 and Ethernet 1 LAN segments is bridged, instead of routed. The **bridge** global configuration command at the end of the example defines the use of the IEEE bridging encapsulation standard for packets bridged on the interfaces joined to this bridging group. The DEC encapsulation technique can also be used by specifying the **dec** keyword at the end of the bridge global configuration command. Bridging is supported in the following instances:

- Ethernet
- Token Ring
- FDDI
- ATM
- X.25 (HDLC)
- Frame Relay (HDLC)
- ISDN



NOTE: When defining the bridge-group command on an active Token Ring or on FDDI interfaces, or when you are adding new interfaces to an existing bridge-group that contains Token Ring or FDDI interfaces, all Token Ring and FDDI interfaces of the bridge-group are immediately reinitialized, disrupting connectivity.


Controlling and Managing Broadcast Packets

Broadcast packets are used by several important Internet protocols. Because broadcast messages have the potential to overload a network, control of these packets is essential to the health of an IP network. Two types of broadcast packets are supported with Cisco IOS software. The first is a directed broadcast.

```
interface Ethernet 1
ip address 10.1.1.2 255.255.255.0
bridge-group 1
!
bridge 1 protocol ieee
```

The value of the bridge-group command can be anywhere from 1 to 63. In the example, IP and all other protocol traffic between Ethernet 0 and Ethernet 1 LAN segments is bridged, instead of routed. The **bridge** global configuration command at the end of the example defines the use of the IEEE bridging encapsulation standard for packets bridged on the interfaces joined to this bridging group. The DEC encapsulation technique can also be used by specifying the **dec** keyword at the end of the bridge global configuration command. Bridging is supported in the following instances:

- Ethernet
- Token Ring
- FDDI
- ATM
- X.25 (HDLC)
- Frame Relay (HDLC)
- ISDN



NOTE: When defining the bridge-group command on an active Token Ring or on FDDI interfaces, or when you are adding new interfaces to an existing bridge-group that contains Token Ring or FDDI interfaces, all Token Ring and FDDI interfaces of the bridge-group are immediately reinitialized, disrupting connectivity.

Controlling and Managing Broadcast Packets

Broadcast packets are used by several important Internet protocols. Because broadcast messages have the potential to overload a network, control of these packets is essential to the health of an IP network. Two types of broadcast packets are supported with Cisco IOS software. The first is a directed broadcast.

A directed broadcast is an IP address that uses the network and or subnet field of the IP address. For example, the 172.16.0.0 network can have a directed broadcast of 172.16.255.255 to which all hosts will respond. Let's assume that the subnet mask, 255.255.255.0, is applied to this Class B address. Applying this mask to the third octet of the IP address is the subnet field. Now the 172.16.30.255 address becomes a broadcast address for all hosts on the 172.16.30.0 subnetwork. Upon recognizing this IP address as a directed broadcast, the router will re-encapsulate the packet in a frame with a broadcast destination MAC address of x'FFFF.FFFF.FFFF', causing all devices to access the frame for processing.

Contrast this to flooding, (also known as local broadcast) the second type of supported broadcast. A flooding broadcast is answered by all devices on all networks. A flooding broadcast is defined by the destination IP address being all ones, or 255.255.255.255 in decimal notation. Sending a broadcast of this nature can cause serious overload on the network, known as a broadcast storm. Routers by default will not forward an all ones IP destination address, thereby keeping the broadcast on the local network.

Enable Directed Broadcast-to-Physical Broadcast Translation

Directed broadcasts are enabled per interface. Prior to Cisco IOS 12.0, directed broadcast translation to a physical broadcast would default as enabled and thereby forward directed broadcasts using a MAC broadcast address. With Release 12.0 of Cisco IOS, this directed broadcast is disabled by default. This is due in large part to the myriad of broadcast storms created by deviants on the Internet using attacks called "smurf" and "fraggle." These attacks place a known IP address within the broadcast as the source of the packet, even though the true source is outside on the Internet. Therefore, in Release 12.0 of Cisco IOS software, directed broadcasts must be explicitly defined on any interface that requires the feature. Let's use the following example to understand how Cisco IOS controls and manages IP broadcasts.

```
ip forward-protocol spanning-tree
 bridge 1 protocol dec
 access-list 201 deny 0x0000 0xFFFF
 interface ethernet 0
 bridge-group 1
 bridge-group 1 input-type-list 201
 interface ethernet 1
 bridge-group 1
 bridge-group 1 input-type-list 201
 interface serial 0
 bridge-group 1
```

```
bridge-group 1 input-type-list 201
interface serial 1
bridge-group 1
bridge-group 1 input-type-list 201
```

In this example, Ethernet 0 and Ethernet 1, along with the two serial interfaces, support the flooding of default protocols TFTP, DNS, Time, NetBIOS, and BOOTP. This example provides the capability to flood a network with the default protocols. The **ip forward-protocol spanning-tree** command enables transparent bridging, which is a requirement for broadcast flooding. The **access-list** filters all protocols from being bridged, with the exception of the default protocols. This access list is applied to each of the interfaces as an input filter, thereby protecting the propagation of all unwanted protocol broadcasts.

Packets must meet the following criteria to be considered for flooding:

- The packet must be a MAC-level broadcast.
- The packet must be an IP-level broadcast.
- The packet must be a TFTP, DNS, Time, NetBIOS, ND, or BOOTP packet, or a UDP protocol specified by the **ip forward-protocol udp** global configuration command.
- The packet's time-to-live (TTL) value must be at least two.

These are the same criteria used to consider packet forwarding using IP helper addresses.

To facilitate network broadcast requirements, such as workstations getting their network configuration parameters using Dynamic Host Configuration Protocol (DHCP), Cisco IOS redirects the broadcast to a specific IP address. Using the following example,

```
ip forward-protocol udp
!
interface ethernet 0
 ip address 10.1.1.1 255.255.255.0
 ip helper-address 10.4.2.7
interface ethernet 1
 ip address 10.1.2.1 255.255.255.0
 ip helper-address 10.4.2.7
```

the **ip forward-protocol udp** global command is entered to allow the following UDP application protocols to be forwarded:

- Trivial File Transfer Protocol (TFTP) (port 69)
- DNS (port 53)
- Time service (port 37)
- NetBIOS Name Server (port 137)
- NetBIOS Datagram Server (port 138)

- Boot Protocol (BOOTP) client and server datagrams (ports 67 and 68)
- TACACS service (port 49)

Other UDP applications' protocol broadcasts can be forwarded by supplying the appropriate UDP port number at the end of the **ip forward-protocol udp** command. In Figure 13-8, we see that the workstations on Ethernet 0 and Ethernet 1 send a BOOTP (DHCP) broadcast to locate a DHCP server for acquiring network configuration parameters. The router forwards the request to the IP address provided by the ip helper-address defined for each interface. Multiple ip helper-address commands can be entered for each interface. Having multiple ip helper-address commands enables load distribution, redundant service, and multiple services supported for each interface.

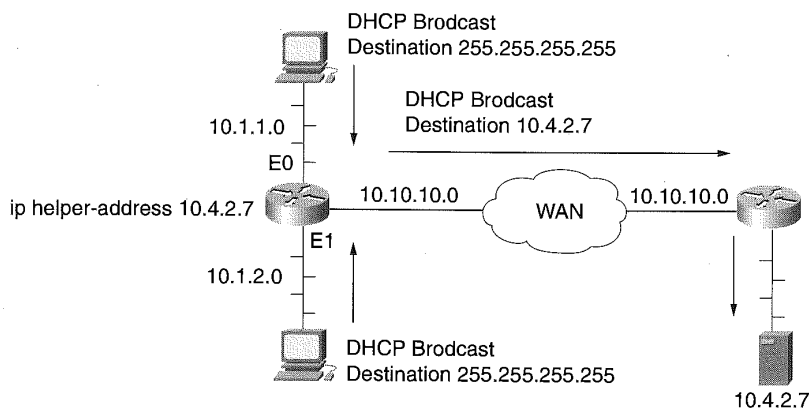
Configure IP Services

Cisco IOS supports many different services specific to IP. Among the various services discussed here are ICMP, filtering with access lists, redundancy using HSRP, accounting, and performance tuning.

Disable ICMP Unreachable Messages

Cisco IOS automatically replies to an ICMP message for an unknown protocol by sending an ICMP Protocol Unreachable message back to the sending IP address. Likewise, if the ICMP message is undeliverable to the destination IP host address, it returns an ICMP Host Unreachable message to the sender.

Figure 13-8
Directing UDP DHCP broadcast using ip helper-address.



Entering the following command for any interface under interface configuration mode disables the ICMP Unreachable messages:

no ip unreachable

Reenabling the ICMP Unreachable message for an interface is accomplished by entering the

ip unreachable

interface configuration command. You may want to disable ICMP Unreachable messages on interfaces connecting to outside networks. Hackers sometimes use this information to obtain the reachability and available protocols that are allowed through the firewall.

Disable ICMP Redirect Messages

By default, the ICMP Redirect message is enabled for all interfaces supporting IP on the router. The ICMP Redirect message is sent to the originator of the packet, indicating the destination of the packet is either on the same subnet or there is a more direct path to the destination. In this case, the originating host places the IP address provided by the router that sent the ICMP Redirect message in the packet. The ICMP Redirect is disabled using the following interface configuration command:

no ip redirects

Again, hackers will use the ICMP Redirect message to break into a private network. To reenable the ICMP Redirect messages, enter the

ip redirects

interface configuration command for the interface on which ICMP redirects are desired.

Disable ICMP Mask Reply Messages

Another technique used by hackers is to request the IP address mask of an IP address by using the ICMP Mask Request message. The ICMP Mask Reply provides the final piece of the puzzle for hackers to gain access to your network. To disable the ICMP Mask Reply, enter the following command under interface configuration mode:

no ip mask-reply

Again, to reenable the command on a specific interface that has already disabled this function, enter the following command under interface configuration mode:

ip mask-reply

Support for Fragmenting Large IP Packets

The maximum transmission unit (MTU) size is applied to all defined Cisco router interfaces. The MTU size is the largest size packet supported on the defined interface. The MTU defaults are different for each interface type defined on the router. The largest possible size MTU on a route when using IP is determined by the IP Path MTU Discovery mechanism.

All Cisco routers employ the IP Path Discovery mechanism. As a router receives a packet on an interface, it determines the size of the packet, the MTU size allowed on the next hop interface, and whether the don't fragment (DF) bit is set by the origin IP host. Figure 13-9 illustrates an example in which the sending device sends a packet larger than that allowed on the destination LAN segment and the sending device has the DF bit set.

Figure 13-9
IP Path MTU
Discovery mechanism
with Cisco routers.

